

# Intel<sup>®</sup> 64 and IA-32 Architectures Software Developer's Manual

Documentation Changes

---

March 2018

**Notice:** The Intel<sup>®</sup> 64 and IA-32 architectures may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are documented in the specification updates.

Document Number: 252046-058



Intel technologies features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at [intel.com](http://intel.com), or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting <http://www.intel.com/design/literature.htm>.

Intel, the Intel logo, Intel Atom, Intel Core, Intel SpeedStep, MMX, Pentium, VTune, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

\*Other names and brands may be claimed as the property of others.

Copyright © 1997-2018, Intel Corporation. All Rights Reserved.



# Contents

---

|                                     |   |
|-------------------------------------|---|
| Revision History . . . . .          | 4 |
| Preface . . . . .                   | 7 |
| Summary Tables of Changes . . . . . | 8 |
| Documentation Changes. . . . .      | 9 |



# Revision History

---

| Revision | Description  | Date           |
|----------|--|----------------|
| -001     | <ul style="list-style-type: none"><li>Initial release</li></ul>  | November 2002  |
| -002     | <ul style="list-style-type: none"><li>Added 1-10 Documentation Changes.</li><li>Removed old Documentation Changes items that already have been incorporated in the published Software Developer's manual</li></ul>                                     | December 2002  |
| -003     | <ul style="list-style-type: none"><li>Added 9 -17 Documentation Changes.</li><li>Removed Documentation Change #6 - References to bits Gen and Len Deleted.</li><li>Removed Documentation Change #4 - VIF Information Added to CLI Discussion</li></ul> | February 2003  |
| -004     | <ul style="list-style-type: none"><li>Removed Documentation changes 1-17.</li><li>Added Documentation changes 1-24.</li></ul>  | June 2003      |
| -005     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-24.</li><li>Added Documentation Changes 1-15.</li></ul>  | September 2003 |
| -006     | <ul style="list-style-type: none"><li>Added Documentation Changes 16- 34.</li></ul>  | November 2003  |
| -007     | <ul style="list-style-type: none"><li>Updated Documentation changes 14, 16, 17, and 28.</li><li>Added Documentation Changes 35-45.</li></ul>   | January 2004   |
| -008     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-45.</li><li>Added Documentation Changes 1-5.</li></ul>   | March 2004     |
| -009     | <ul style="list-style-type: none"><li>Added Documentation Changes 7-27.</li></ul>  | May 2004       |
| -010     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-27.</li><li>Added Documentation Changes 1.</li></ul>   | August 2004    |
| -011     | <ul style="list-style-type: none"><li>Added Documentation Changes 2-28.</li></ul>  | November 2004  |
| -012     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-28.</li><li>Added Documentation Changes 1-16.</li></ul>  | March 2005     |
| -013     | <ul style="list-style-type: none"><li>Updated title.</li><li>There are no Documentation Changes for this revision of the document.</li></ul>   | July 2005      |
| -014     | <ul style="list-style-type: none"><li>Added Documentation Changes 1-21.</li></ul>  | September 2005 |
| -015     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-21.</li><li>Added Documentation Changes 1-20.</li></ul>  | March 9, 2006  |
| -016     | <ul style="list-style-type: none"><li>Added Documentation changes 21-23.</li></ul>   | March 27, 2006 |
| -017     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-23.</li><li>Added Documentation Changes 1-36.</li></ul>  | September 2006 |
| -018     | <ul style="list-style-type: none"><li>Added Documentation Changes 37-42.</li></ul>   | October 2006   |
| -019     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-42.</li><li>Added Documentation Changes 1-19.</li></ul>  | March 2007     |
| -020     | <ul style="list-style-type: none"><li>Added Documentation Changes 20-27.</li></ul>   | May 2007       |
| -021     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-27.</li><li>Added Documentation Changes 1-6</li></ul>  | November 2007  |
| -022     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-6</li><li>Added Documentation Changes 1-6</li></ul>  | August 2008    |
| -023     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-6</li><li>Added Documentation Changes 1-21</li></ul>   | March 2009     |



| Revision | Description  | Date           |
|----------|--|----------------|
| -024     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-21</li> <li>Added Documentation Changes 1-16</li> </ul> | June 2009      |
| -025     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-16</li> <li>Added Documentation Changes 1-18</li> </ul> | September 2009 |
| -026     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-18</li> <li>Added Documentation Changes 1-15</li> </ul> | December 2009  |
| -027     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-15</li> <li>Added Documentation Changes 1-24</li> </ul> | March 2010     |
| -028     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-24</li> <li>Added Documentation Changes 1-29</li> </ul> | June 2010      |
| -029     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-29</li> <li>Added Documentation Changes 1-29</li> </ul> | September 2010 |
| -030     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-29</li> <li>Added Documentation Changes 1-29</li> </ul> | January 2011   |
| -031     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-29</li> <li>Added Documentation Changes 1-29</li> </ul> | April 2011     |
| -032     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-29</li> <li>Added Documentation Changes 1-14</li> </ul> | May 2011       |
| -033     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-14</li> <li>Added Documentation Changes 1-38</li> </ul> | October 2011   |
| -034     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-38</li> <li>Added Documentation Changes 1-16</li> </ul> | December 2011  |
| -035     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-16</li> <li>Added Documentation Changes 1-18</li> </ul> | March 2012     |
| -036     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-18</li> <li>Added Documentation Changes 1-17</li> </ul> | May 2012       |
| -037     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-17</li> <li>Added Documentation Changes 1-28</li> </ul> | August 2012    |
| -038     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-28</li> <li>Add Documentation Changes 1-22</li> </ul>   | January 2013   |
| -039     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-22</li> <li>Add Documentation Changes 1-17</li> </ul>   | June 2013      |
| -040     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-17</li> <li>Add Documentation Changes 1-24</li> </ul>   | September 2013 |
| -041     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-24</li> <li>Add Documentation Changes 1-20</li> </ul>   | February 2014  |
| -042     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-20</li> <li>Add Documentation Changes 1-8</li> </ul>    | February 2014  |
| -043     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-8</li> <li>Add Documentation Changes 1-43</li> </ul>    | June 2014      |
| -044     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-43</li> <li>Add Documentation Changes 1-12</li> </ul>   | September 2014 |
| -045     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-12</li> <li>Add Documentation Changes 1-22</li> </ul>   | January 2015   |
| -046     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-22</li> <li>Add Documentation Changes 1-25</li> </ul>   | April 2015     |
| -047     | <ul style="list-style-type: none"> <li>Removed Documentation Changes 1-25</li> <li>Add Documentation Changes 1-19</li> </ul>   | June 2015      |



| Revision | Description   | Date           |
|----------|---|----------------|
| -048     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-19</li><li>Add Documentation Changes 1-33</li></ul> | September 2015 |
| -049     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-33</li><li>Add Documentation Changes 1-33</li></ul> | December 2015  |
| -050     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-33</li><li>Add Documentation Changes 1-9</li></ul>  | April 2016     |
| -051     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-9</li><li>Add Documentation Changes 1-20</li></ul>  | June 2016      |
| -052     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-20</li><li>Add Documentation Changes 1-22</li></ul> | September 2016 |
| -053     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-22</li><li>Add Documentation Changes 1-26</li></ul> | December 2016  |
| -054     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-26</li><li>Add Documentation Changes 1-20</li></ul> | March 2017     |
| -055     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-20</li><li>Add Documentation Changes 1-28</li></ul> | July 2017      |
| -056     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-28</li><li>Add Documentation Changes 1-18</li></ul> | October 2017   |
| -057     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-18</li><li>Add Documentation Changes 1-29</li></ul> | December 2017  |
| -058     | <ul style="list-style-type: none"><li>Removed Documentation Changes 1-29</li><li>Add Documentation Changes 1-17</li></ul> | March 2018     |

§

# Preface

---

This document is an update to the specifications contained in the [Affected Documents](#) table below. This document is a compilation of device and documentation errata, specification clarifications and changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools.

## Affected Documents

| Document Title  | Document Number/<br>Location |
|---|------------------------------|
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture</i>                | 253665                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference, A-L</i>   | 253666                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference, M-U</i>   | 253667                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C: Instruction Set Reference, V-Z</i>   | 326018                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2D: Instruction Set Reference</i>        | 334569                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide, Part 1</i> | 253668                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2</i> | 253669                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C: System Programming Guide, Part 3</i> | 326019                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3D: System Programming Guide, Part 4</i> | 332831                       |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model Specific Registers</i>          | 335592                       |

## Nomenclature

**Documentation Changes** include typos, errors, or omissions from the current published specifications. These will be incorporated in any new release of the specification.

# Summary Tables of Changes

---

The following table indicates documentation changes which apply to the Intel® 64 and IA-32 architectures. This table uses the following notations:

## Codes Used in Summary Tables

Change bar to left of table row indicates this erratum is either new or modified from the previous version of the document.

## Documentation Changes

| No. | DOCUMENTATION CHANGES            |
|-----|----------------------------------|
| 1   | Updates to Chapter 1, Volume 1   |
| 2   | Updates to Chapter 13, Volume 1  |
| 3   | Updates to Chapter 1, Volume 2A  |
| 4   | Updates to Chapter 3, Volume 2A  |
| 5   | Updates to Chapter 4, Volume 2B  |
| 6   | Updates to Chapter 5, Volume 2C  |
| 7   | Updates to Chapter 7, Volume 2D  |
| 8   | Updates to Chapter 1, Volume 3A  |
| 9   | Updates to Chapter 4, Volume 3A  |
| 10  | Updates to Chapter 14, Volume 3B |
| 11  | Updates to Chapter 18, Volume 3B |
| 12  | Updates to Chapter 19, Volume 3B |
| 13  | Updates to Chapter 24, Volume 3B |
| 14  | Updates to Chapter 35, Volume 3C |
| 15  | Updates to Chapter 40, Volume 3D |
| 16  | Updates to Chapter 1, Volume 4   |
| 17  | Updates to Chapter 2, Volume 4   |



# *Documentation Changes*

---

Changes to the Intel® 64 and IA-32 Architectures Software Developer's Manual volumes follow, and are listed by chapter. Only chapters with changes are included in this document.



## 1. Updates to Chapter 1, Volume 1

Change bars show changes to Chapter 1 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*.

-----  
Change to this chapter: Updates to processors covered by manual; added 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series.

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture* (order number 253665) is part of a set that describes the architecture and programming environment of Intel® 64 and IA-32 architecture processors. Other volumes in this set are:

- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D: Instruction Set Reference* (order numbers 253666, 253667, 326018 and 334569).
- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D: System Programming Guide* (order numbers 253668, 253669, 326019 and 332831).
- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers* (order number 335592).

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, describes the basic architecture and programming environment of Intel 64 and IA-32 processors. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D*, describe the instruction set of the processor and the opcode structure. These volumes apply to application programmers and to programmers who write operating systems or executives. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D*, describe the operating-system support environment of Intel 64 and IA-32 processors. These volumes target operating-system and BIOS designers. In addition, the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B*, addresses the programming environment for classes of software that host operating systems. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*, describes the model-specific registers of Intel 64 and IA-32 processors.

## 1.1 INTEL® 64 AND IA-32 PROCESSORS COVERED IN THIS MANUAL

This manual set includes information pertaining primarily to the most recent Intel 64 and IA-32 processors, which include:

- Pentium® processors
- P6 family processors
- Pentium® 4 processors
- Pentium® M processors
- Intel® Xeon® processors
- Pentium® D processors
- Pentium® processor Extreme Editions
- 64-bit Intel® Xeon® processors
- Intel® Core™ Duo processor
- Intel® Core™ Solo processor
- Dual-Core Intel® Xeon® processor LV
- Intel® Core™2 Duo processor
- Intel® Core™2 Quad processor Q6000 series
- Intel® Xeon® processor 3000, 3200 series
- Intel® Xeon® processor 5000 series
- Intel® Xeon® processor 5100, 5300 series
- Intel® Core™2 Extreme processor X7000 and X6800 series
- Intel® Core™2 Extreme processor QX6000 series
- Intel® Xeon® processor 7100 series

## ABOUT THIS MANUAL

- Intel® Pentium® Dual-Core processor
- Intel® Xeon® processor 7200, 7300 series
- Intel® Xeon® processor 5200, 5400, 7400 series
- Intel® Core™2 Extreme processor QX9000 and X9000 series
- Intel® Core™2 Quad processor Q9000 series
- Intel® Core™2 Duo processor E8000, T9000 series
- Intel® Atom™ processor family
- Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are built from 45 nm and 32 nm processes
- Intel® Core™ i7 processor
- Intel® Core™ i5 processor
- Intel® Xeon® processor E7-8800/4800/2800 product families
- Intel® Core™ i7-3930K processor
- 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series
- Intel® Xeon® processor E3-1200 product family
- Intel® Xeon® processor E5-2400/1400 product family
- Intel® Xeon® processor E5-4600/2600/1600 product family
- 3rd generation Intel® Core™ processors
- Intel® Xeon® processor E3-1200 v2 product family
- Intel® Xeon® processor E5-2400/1400 v2 product families
- Intel® Xeon® processor E5-4600/2600/1600 v2 product families
- Intel® Xeon® processor E7-8800/4800/2800 v2 product families
- 4th generation Intel® Core™ processors
- The Intel® Core™ M processor family
- Intel® Core™ i7-59xx Processor Extreme Edition
- Intel® Core™ i7-49xx Processor Extreme Edition
- Intel® Xeon® processor E3-1200 v3 product family
- Intel® Xeon® processor E5-2600/1600 v3 product families
- 5th generation Intel® Core™ processors
- Intel® Xeon® processor D-1500 product family
- Intel® Xeon® processor E5 v4 family
- Intel® Atom™ processor X7-Z8000 and X5-Z8000 series
- Intel® Atom™ processor Z3400 series
- Intel® Atom™ processor Z3500 series
- 6th generation Intel® Core™ processors
- Intel® Xeon® processor E3-1500m v5 product family
- 7th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series
- Intel® Xeon® Processor Scalable Family
- 8th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series

P6 family processors are IA-32 processors based on the P6 family microarchitecture. This includes the Pentium® Pro, Pentium® II, Pentium® III, and Pentium® III Xeon® processors.

The Pentium® 4, Pentium® D, and Pentium® processor Extreme Editions are based on the Intel NetBurst® microarchitecture. Most early Intel® Xeon® processors are based on the Intel NetBurst® microarchitecture. Intel Xeon processor 5000, 7100 series are based on the Intel NetBurst® microarchitecture.

The Intel® Core™ Duo, Intel® Core™ Solo and dual-core Intel® Xeon® processor LV are based on an improved Pentium® M processor microarchitecture.

The Intel® Xeon® processor 3000, 3200, 5100, 5300, 7200, and 7300 series, Intel® Pentium® dual-core, Intel® Core™2 Duo, Intel® Core™2 Quad, and Intel® Core™2 Extreme processors are based on Intel® Core™ microarchitecture.

The Intel® Xeon® processor 5200, 5400, 7400 series, Intel® Core™2 Quad processor Q9000 series, and Intel® Core™2 Extreme processors QX9000, X9000 series, Intel® Core™2 processor E8000 series are based on Enhanced Intel® Core™ microarchitecture.

The Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are based on the Intel® Atom™ microarchitecture and supports Intel 64 architecture.

P6 family, Pentium® M, Intel® Core™ Solo, Intel® Core™ Duo processors, dual-core Intel® Xeon® processor LV, and early generations of Pentium 4 and Intel Xeon processors support IA-32 architecture. The Intel® Atom™ processor Z5xx series support IA-32 architecture.

The Intel® Xeon® processor 3000, 3200, 5000, 5100, 5200, 5300, 5400, 7100, 7200, 7300, 7400 series, Intel® Core™2 Duo, Intel® Core™2 Extreme, Intel® Core™2 Quad processors, Pentium® D processors, Pentium® Dual-Core processor, newer generations of Pentium 4 and Intel Xeon processor family support Intel® 64 architecture.

The Intel® Core™ i7 processor and Intel® Xeon® processor 3400, 5500, 7500 series are based on 45 nm Intel® microarchitecture code name Nehalem. Intel® microarchitecture code name Westmere is a 32 nm version of Intel® microarchitecture code name Nehalem. Intel® Xeon® processor 5600 series, Intel Xeon processor E7 and various Intel Core i7, i5, i3 processors are based on Intel® microarchitecture code name Westmere. These processors support Intel 64 architecture.

The Intel® Xeon® processor E5 family, Intel® Xeon® processor E3-1200 family, Intel® Xeon® processor E7-8800/4800/2800 product families, Intel® Core™ i7-3930K processor, and 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series are based on the Intel® microarchitecture code name Sandy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E7-8800/4800/2800 v2 product families, Intel® Xeon® processor E3-1200 v2 product family and 3rd generation Intel® Core™ processors are based on the Intel® microarchitecture code name Ivy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E5-4600/2600/1600 v2 product families, Intel® Xeon® processor E5-2400/1400 v2 product families and Intel® Core™ i7-49xx Processor Extreme Edition are based on the Intel® microarchitecture code name Ivy Bridge-E and support Intel 64 architecture.

The Intel® Xeon® processor E3-1200 v3 product family and 4th Generation Intel® Core™ processors are based on the Intel® microarchitecture code name Haswell and support Intel 64 architecture.

The Intel® Xeon® processor E5-2600/1600 v3 product families and the Intel® Core™ i7-59xx Processor Extreme Edition are based on the Intel® microarchitecture code name Haswell-E and support Intel 64 architecture.

The Intel® Atom™ processor Z8000 series is based on the Intel microarchitecture code name Airmont.

The Intel® Atom™ processor Z3400 series and the Intel® Atom™ processor Z3500 series are based on the Intel microarchitecture code name Silvermont.

The Intel® Core™ M processor family, 5th generation Intel® Core™ processors, Intel® Xeon® processor D-1500 product family and the Intel® Xeon® processor E5 v4 family are based on the Intel® microarchitecture code name Broadwell and support Intel 64 architecture.

The Intel® Xeon® Processor Scalable Family, Intel® Xeon® processor E3-1500m v5 product family and 6th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Skylake and support Intel 64 architecture.

The 7th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Kaby Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series is based on the Intel® microarchitecture code name Knights Landing and supports Intel 64 architecture.

The 8th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Coffee Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series is based on the Intel® microarchitecture code name Knights Mill and supports Intel 64 architecture.

IA-32 architecture is the instruction set architecture and programming environment for Intel's 32-bit microprocessors. Intel® 64 architecture is the instruction set architecture and programming environment which is the superset of Intel's 32-bit and 64-bit architectures. It is compatible with the IA-32 architecture.

## 1.2 OVERVIEW OF VOLUME 1: BASIC ARCHITECTURE

A description of this manual's content follows:

**Chapter 1 — About This Manual.** Gives an overview of all five volumes of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*. It also describes the notational conventions in these manuals and lists related Intel manuals and documentation of interest to programmers and hardware designers.

**Chapter 2 — Intel® 64 and IA-32 Architectures.** Introduces the Intel 64 and IA-32 architectures along with the families of Intel processors that are based on these architectures. It also gives an overview of the common features found in these processors and brief history of the Intel 64 and IA-32 architectures.

**Chapter 3 — Basic Execution Environment.** Introduces the models of memory organization and describes the register set used by applications.

**Chapter 4 — Data Types.** Describes the data types and addressing modes recognized by the processor; provides an overview of real numbers and floating-point formats and of floating-point exceptions.

**Chapter 5 — Instruction Set Summary.** Lists all Intel 64 and IA-32 instructions, divided into technology groups.

**Chapter 6 — Procedure Calls, Interrupts, and Exceptions.** Describes the procedure stack and mechanisms provided for making procedure calls and for servicing interrupts and exceptions.

**Chapter 7 — Programming with General-Purpose Instructions.** Describes basic load and store, program control, arithmetic, and string instructions that operate on basic data types, general-purpose and segment registers; also describes system instructions that are executed in protected mode.

**Chapter 8 — Programming with the x87 FPU.** Describes the x87 floating-point unit (FPU), including floating-point registers and data types; gives an overview of the floating-point instruction set and describes the processor's floating-point exception conditions.

**Chapter 9 — Programming with Intel® MMX™ Technology.** Describes Intel MMX technology, including MMX registers and data types; also provides an overview of the MMX instruction set.

**Chapter 10 — Programming with Intel® Streaming SIMD Extensions (Intel® SSE).** Describes SSE extensions, including XMM registers, the MXCSR register, and packed single-precision floating-point data types; provides an overview of the SSE instruction set and gives guidelines for writing code that accesses the SSE extensions.

**Chapter 11 — Programming with Intel® Streaming SIMD Extensions 2 (Intel® SSE2).** Describes SSE2 extensions, including XMM registers and packed double-precision floating-point data types; provides an overview of the SSE2 instruction set and gives guidelines for writing code that accesses SSE2 extensions. This chapter also describes SIMD floating-point exceptions that can be generated with SSE and SSE2 instructions. It also provides general guidelines for incorporating support for SSE and SSE2 extensions into operating system and applications code.

**Chapter 12 — Programming with Intel® Streaming SIMD Extensions 3 (Intel® SSE3), Supplemental Streaming SIMD Extensions 3 (SSSE3), Intel® Streaming SIMD Extensions 4 (Intel® SSE4) and Intel® AES New Instructions (Intel® AESNI).** Provides an overview of the SSE3 instruction set, Supplemental SSE3, SSE4, AESNI instructions, and guidelines for writing code that accesses these extensions.

**Chapter 13 — Managing State Using the XSAVE Feature Set.** Describes the XSAVE feature set instructions and explains how software can enable the XSAVE feature set and XSAVE-enabled features.

**Chapter 14 — Programming with AVX, FMA and AVX2.** Provides an overview of the Intel® AVX instruction set, FMA and Intel AVX2 extensions and gives guidelines for writing code that accesses these extensions.

**Chapter 15 — Programming with Intel Transactional Synchronization Extensions.** Describes the instruction extensions that support lock elision techniques to improve the performance of multi-threaded software with contended locks.

**Chapter 16 — Input/Output.** Describes the processor’s I/O mechanism, including I/O port addressing, I/O instructions, and I/O protection mechanisms.

**Chapter 17 — Processor Identification and Feature Determination.** Describes how to determine the CPU type and features available in the processor.

**Appendix A — EFLAGS Cross-Reference.** Summarizes how the IA-32 instructions affect the flags in the EFLAGS register.

**Appendix B — EFLAGS Condition Codes.** Summarizes how conditional jump, move, and ‘byte set on condition code’ instructions use condition code flags (OF, CF, ZF, SF, and PF) in the EFLAGS register.

**Appendix C — Floating-Point Exceptions Summary.** Summarizes exceptions raised by the x87 FPU floating-point and SSE/SSE2/SSE3 floating-point instructions.

**Appendix D — Guidelines for Writing x87 FPU Exception Handlers.** Describes how to design and write MS-DOS\* compatible exception handling facilities for FPU exceptions (includes software and hardware requirements and assembly-language code examples). This appendix also describes general techniques for writing robust FPU exception handlers.

**Appendix E — Guidelines for Writing SIMD Floating-Point Exception Handlers.** Gives guidelines for writing exception handlers for exceptions generated by SSE/SSE2/SSE3 floating-point instructions.

## 1.3 NOTATIONAL CONVENTIONS

This manual uses specific notation for data-structure formats, for symbolic representation of instructions, and for hexadecimal and binary numbers. This notation is described below.

### 1.3.1 Bit and Byte Order

In illustrations of data structures in memory, smaller addresses appear toward the bottom of the figure; addresses increase toward the top. Bit positions are numbered from right to left. The numerical value of a set bit is equal to two raised to the power of the bit position. Intel 64 and IA-32 processors are “little endian” machines; this means the bytes of a word are numbered starting from the least significant byte. See Figure 1-1.




Figure 1-1. Bit and Byte Order



## 1.3.2 Reserved Bits and Software Compatibility

In many register and memory layout descriptions, certain bits are marked as reserved. When bits are marked as reserved, it is essential for compatibility with future processors that software treat these bits as having a future, though unknown, effect. The behavior of reserved bits should be regarded as not only undefined, but unpredictable.

Software should follow these guidelines in dealing with reserved bits:

- Do not depend on the states of any reserved bits when testing the values of registers that contain such bits. Mask out the reserved bits before testing.
- Do not depend on the states of any reserved bits when storing to memory or to a register.
- Do not depend on the ability to retain information written into any reserved bits.
- When loading a register, always load the reserved bits with the values indicated in the documentation, if any, or reload them with values previously read from the same register.

### NOTE

Avoid any software dependence upon the state of reserved bits in Intel 64 and IA-32 registers. Depending upon the values of reserved register bits will make software dependent upon the unspecified manner in which the processor handles these bits. Programs that depend upon reserved values risk incompatibility with future processors.

### 1.3.2.1 Instruction Operands

When instructions are represented symbolically, a subset of the IA-32 assembly language is used. In this subset, an instruction has the following format:

```
label: mnemonic argument1, argument2, argument3
```

where:

- A **label** is an identifier which is followed by a colon.
- A **mnemonic** is a reserved name for a class of instruction opcodes which have the same function.
- The operands **argument1**, **argument2**, and **argument3** are optional. There may be from zero to three operands, depending on the opcode. When present, they take the form of either literals or identifiers for data items. Operand identifiers are either reserved names of registers or are assumed to be assigned to data items declared in another part of the program (which may not be shown in the example).

When two operands are present in an arithmetic or logical instruction, the right operand is the source and the left operand is the destination.

For example:

```
LOADREG: MOV EAX, SUBTOTAL
```

In this example, **LOADREG** is a label, **MOV** is the mnemonic identifier of an opcode, **EAX** is the destination operand, and **SUBTOTAL** is the source operand. Some assembly languages put the source and destination in reverse order.

## 1.3.3 Hexadecimal and Binary Numbers

Base 16 (hexadecimal) numbers are represented by a string of hexadecimal digits followed by the character H (for example, 0F82EH). A hexadecimal digit is a character from the following set: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, and F.

Base 2 (binary) numbers are represented by a string of 1s and 0s, sometimes followed by the character B (for example, 1010B). The "B" designation is only used in situations where confusion as to the type of number might arise.

### 1.3.4 Segmented Addressing

The processor uses byte addressing. This means memory is organized and accessed as a sequence of bytes. Whether one or more bytes are being accessed, a byte address is used to locate the byte or bytes memory. The range of memory that can be addressed is called an **address space**.

The processor also supports segmented addressing. This is a form of addressing where a program may have many independent address spaces, called **segments**. For example, a program can keep its code (instructions) and stack in separate segments. Code addresses would always refer to the code space, and stack addresses would always refer to the stack space. The following notation is used to specify a byte address within a segment:

Segment-register:Byte-address

For example, the following segment address identifies the byte at address FF79H in the segment pointed by the DS register:

DS:FF79H

The following segment address identifies an instruction address in the code segment. The CS register points to the code segment and the EIP register contains the address of the instruction.

CS:EIP

### 1.3.5 A New Syntax for CPUID, CR, and MSR Values

Obtain feature flags, status, and system information by using the CPUID instruction, by checking control register bits, and by reading model-specific registers. We are moving toward a new syntax to represent this information. See Figure 1-2.




Figure 1-2. Syntax for CPUID, CR, and MSR Data Presentation

## 1.3.6 Exceptions

An exception is an event that typically occurs when an instruction causes an error. For example, an attempt to divide by zero generates an exception. However, some exceptions, such as breakpoints, occur under other conditions. Some types of exceptions may provide error codes. An error code reports additional information about the error. An example of the notation used to show an exception and error code is shown below:

```
#PF(fault code)
```

This example refers to a page-fault exception under conditions where an error code naming a type of fault is reported. Under some conditions, exceptions that produce error codes may not be able to report an accurate code. In this case, the error code is zero, as shown below for a general-protection exception:

```
#GP(0)
```

## 1.4 RELATED LITERATURE

Literature related to Intel 64 and IA-32 processors is listed and viewable on-line at:

<http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>

See also:

- The data sheet for a particular Intel 64 or IA-32 processor
- The specification update for a particular Intel 64 or IA-32 processor
- Intel® C++ Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Fortran Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Software Development Tools:  
<https://software.intel.com/en-us/intel-sdp-home>
- Intel® 64 and IA-32 Architectures Software Developer's Manual (in one, four or ten volumes):  
<https://software.intel.com/en-us/articles/intel-sdm>
- Intel® 64 and IA-32 Architectures Optimization Reference Manual:  
<https://software.intel.com/en-us/articles/intel-sdm#optimization>
- Intel 64 Architecture x2APIC Specification:  
<http://www.intel.com/content/www/us/en/architecture-and-technology/64-architecture-x2apic-specification.html>
- Intel® Trusted Execution Technology Measured Launched Environment Programming Guide:  
<http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>
- Developing Multi-threaded Applications: A Platform Consistent Approach:  
<https://software.intel.com/sites/default/files/article/147714/51534-developing-multithreaded-applications.pdf>
- Using Spin-Loops on Intel® Pentium® 4 Processor and Intel® Xeon® Processor:  
<https://software.intel.com/sites/default/files/22/30/25602>
- Performance Monitoring Unit Sharing Guide  
<http://software.intel.com/file/30388>

Literature related to selected features in future Intel processors are available at:

- Intel® Architecture Instruction Set Extensions Programming Reference  
<https://software.intel.com/en-us/isa-extensions>
- Intel® Software Guard Extensions (Intel® SGX) Programming Reference  
<https://software.intel.com/en-us/isa-extensions/intel-sgx>

## ABOUT THIS MANUAL

More relevant links are:

- Intel® Developer Zone:  
<https://software.intel.com/en-us>
- Developer centers:  
<http://www.intel.com/content/www/us/en/hardware-developers/developer-centers.html>
- Processor support general link:  
<http://www.intel.com/support/processors/>
- Intel® Hyper-Threading Technology (Intel® HT Technology):  
<http://www.intel.com/technology/platform-technology/hyper-threading/index.htm>

## 2. Updates to Chapter 13, Volume 1

Change bars show changes to Chapter 13 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*.

-----  
Change to this chapter: Updates to XSAVE related to hardware duty cycling.

# CHAPTER 13

## MANAGING STATE USING THE XSAVE FEATURE SET

---

The XSAVE feature set extends the functionality of the FXSAVE and FXRSTOR instructions (see Section 10.5, “FXSAVE and FXRSTOR Instructions”) by supporting the saving and restoring of processor state in addition to the x87 execution environment (**x87 state**) and the registers used by the streaming SIMD extensions (**SSE state**).

The **XSAVE feature set** comprises eight instructions. XGETBV and XSETBV allow software to read and write the extended control register XCR0, which controls the operation of the XSAVE feature set. XSAVE, XSAVEOPT, XSAVEC, and XSAVES are four instructions that save processor state to memory; XRSTOR and XRSTORS are corresponding instructions that load processor state from memory. XGETBV, XSAVE, XSAVEOPT, XSAVEC, and XRSTOR can be executed at any privilege level; XSETBV, XSAVES, and XRSTORS can be executed only if CPL = 0. In addition to XCR0, the XSAVES and XRSTORS instructions are controlled also by the IA32\_XSS MSR (index DA0H).

The XSAVE feature set organizes the state that manages into **state components**. Operation of the instructions is based on **state-component bitmaps** that have the same format as XCR0 and as the IA32\_XSS MSR: each bit corresponds to a state component. Section 13.1 discusses these state components and bitmaps in more detail.

Section 13.2 describes how the processor enumerates support for the XSAVE feature set and for **XSAVE-enabled features** (those features that require use of the XSAVE feature set for their enabling). Section 13.3 explains how software can enable the XSAVE feature set and XSAVE-enabled features.

The XSAVE feature set allows saving and loading processor state from a region of memory called an **XSAVE area**. Section 13.4 presents details of the XSAVE area and its organization. Each XSAVE-managed state component is associated with a section of the XSAVE area. Section 13.5 describes in detail each of the XSAVE-managed state components.

Section 13.7 through Section 13.12 describe the operation of XSAVE, XRSTOR, XSAVEOPT, XSAVEC, XSAVES, and XRSTORS, respectively.

### 13.1 XSAVE-SUPPORTED FEATURES AND STATE-COMPONENT BITMAPS

The XSAVE feature set supports the saving and restoring of **state components**, each of which is a discrete set of processor registers (or parts of registers). In general, each such state component corresponds to a particular CPU feature. Such a feature is **XSAVE-supported**. Some XSAVE-supported features use registers in multiple XSAVE-managed state components.

The XSAVE feature set organizes the state components of the XSAVE-supported features using **state-component bitmaps**. A state-component bitmap comprises 64 bits; each bit in such a bitmap corresponds to a single state component. The following bits are defined in state-component bitmaps:

- Bit 0 corresponds to the state component used for the x87 FPU execution environment (**x87 state**). See Section 13.5.1.
- Bit 1 corresponds to the state component used for registers used by the streaming SIMD extensions (**SSE state**). See Section 13.5.2.
- Bit 2 corresponds to the state component used for the additional register state used by the Intel® Advanced Vector Extensions (**AVX state**). See Section 13.5.3.
- Bits 4:3 correspond to the two state components used for the additional register state used by Intel® Memory Protection Extensions (**MPX state**):
  - State component 3 is used for the 4 128-bit bounds registers BND0–BND3 (**BNDREGS state**).
  - State component 4 is used for the 64-bit user-mode MPX configuration register BNDCFGU and the 64-bit MPX status register BNDSTATUS (**BNDCSR state**).
- Bits 7:5 correspond to the three state components used for the additional register state used by Intel® Advanced Vector Extensions 512 (**AVX-512 state**):
  - State component 5 is used for the 8 64-bit opmask registers k0–k7 (**opmask state**).

- State component 6 is used for the upper 256 bits of the registers ZMM0–ZMM15. These 16 256-bit values are denoted ZMM0\_H–ZMM15\_H (ZMM\_Hi256 state).
- State component 7 is used for the 16 512-bit registers ZMM16–ZMM31 (Hi16\_ZMM state).
- Bit 8 corresponds to the state component used for the Intel Processor Trace MSRs (PT state).
- Bit 9 corresponds to the state component used for the protection-key feature’s register PKRU (PKRU state). See Section 13.5.7.
- Bit 13 corresponds to the state component used for an MSR used to control hardware duty cycling (HDC state). See Section 13.5.8.

Bits in the ranges 62:14 and 12:10 are not currently defined in state-component bitmaps and are reserved for future expansion. As individual state component is defined within bits 62:11, additional sub-sections are updated within Section 13.5 over time. Bit 63 is used for special functionality in some bitmaps and does not correspond to any state component.

The state component corresponding to bit *i* of state-component bitmaps is called **state component *i***. Thus, x87 state is state component 0; SSE state is state component 1; AVX state is state component 2; MPX state comprises state components 3–4; AVX-512 state comprises state components 5–7; PT state is state component 8; PKRU state is state component 9; and HDC state is state component 13.

The XSAVE feature set uses state-component bitmaps in multiple ways. Most of the instructions use an implicit operand (in EDX:EAX), called the **instruction mask**, which is the state-component bitmap that specifies the state components on which the instruction operates.

Some state components are **user state components**, and they can be managed by the entire XSAVE feature set. Other state components are **supervisor state components**, and they can be managed only by XSAVES and XRSTORS. The state components corresponding to bit 9 and to bits in the range 7:0 are user state components, PT state (corresponding to bit 8) and HDC state (corresponding to bit 13) are supervisor state components.

Extended control register XCR0 contains a state-component bitmap that specifies the user state components that software has enabled the XSAVE feature set to manage. If the bit corresponding to a state component is clear in XCR0, instructions in the XSAVE feature set will not operate on that state component, regardless of the value of the instruction mask.

The IA32\_XSS MSR (index DA0H) contains a state-component bitmap that specifies the supervisor state components that software has enabled XSAVES and XRSTORS to manage (XSAVE, XSAVEC, XSAVEOPT, and XRSTOR cannot manage supervisor state components). If the bit corresponding to a state component is clear in the IA32\_XSS MSR, XSAVES and XRSTORS will not operate on that state component, regardless of the value of the instruction mask.

Some XSAVE-supported features can be used only if XCR0 has been configured so that the features’ state components can be managed by the XSAVE feature set. (This applies only to features with user state components.) Such state components and features are **XSAVE-enabled**. In general, the processor will not modify (or allow modification of) the registers of a state component of an XSAVE-enabled feature if the bit corresponding to that state component is clear in XCR0. (If software clears such a bit in XCR0, the processor preserves the corresponding state component.) If an XSAVE-enabled feature has not been fully enabled in XCR0, execution of any instruction defined for that feature causes an invalid-opcode exception (#UD).

As will be explained in Section 13.3, the XSAVE feature set is enabled only if CR4.OSXSAVE[bit 18] = 1. If CR4.OSXSAVE = 0, the processor treats XSAVE-enabled state features and their state components as if all bits in XCR0 were clear; the state components cannot be modified and the features’ instructions cannot be executed.

The state components for x87 state, for SSE state, for PT state, for PKRU state, and for HDC state are XSAVE-managed but the corresponding features are not XSAVE-enabled. Processors allow modification of this state, as well as execution of x87 FPU instructions and SSE instructions and use of Intel Processor Trace, protection keys, and hardware duty cycling regardless of the value of CR4.OSXSAVE and XCR0.



## 13.2 ENUMERATION OF CPU SUPPORT FOR XSAVE INSTRUCTIONS AND XSAVE-SUPPORTED FEATURES

A processor enumerates support for the XSAVE feature set and for features supported by that feature set using the CPUID instruction. The following items provide specific details:

- CPUID.1:ECX.XSAVE[bit 26] enumerates general support for the XSAVE feature set:
  - If this bit is 0, the processor does not support any of the following instructions: XGETBV, XRSTOR, XRSTORS, XSAVE, XSAVEC, XSAVEOPT, XSAVES, and XSETBV; the processor provides no further enumeration through CPUID function 0DH (see below).
  - If this bit is 1, the processor supports the following instructions: XGETBV, XRSTOR, XSAVE, and XSETBV.<sup>1</sup> Further enumeration is provided through CPUID function 0DH.
- CP4.OSXSAVE can be set to 1 if and only if CPUID.1:ECX.XSAVE[bit 26] is enumerated as 1.
- CPUID function 0DH enumerates details of CPU support through a set of sub-functions. Software selects a specific sub-function by the value placed in the ECX register. The following items provide specific details:
  - CPUID function 0DH, sub-function 0.
    - EDX:EAX is a bitmap of all the user state components that can be managed using the XSAVE feature set. A bit can be set in XCR0 if and only if the corresponding bit is set in this bitmap. Every processor that supports the XSAVE feature set will set EAX[0] (x87 state) and EAX[1] (SSE state).  
If  $EAX[i] = 1$  (for  $1 < i < 32$ ) or  $EDX[i-32] = 1$  (for  $32 \leq i < 63$ ), sub-function  $i$  enumerates details for state component  $i$  (see below).
    - ECX enumerates the size (in bytes) required by the XSAVE instruction for an XSAVE area containing all the user state components supported by this processor.
    - EBX enumerates the size (in bytes) required by the XSAVE instruction for an XSAVE area containing all the user state components corresponding to bits currently set in XCR0.
  - CPUID function 0DH, sub-function 1.
    - EAX[0] enumerates support for the XSAVEOPT instruction. The instruction is supported if and only if this bit is 1. If  $EAX[0] = 0$ , execution of XSAVEOPT causes an invalid-opcode exception (#UD).
    - EAX[1] enumerates support for **compaction extensions** to the XSAVE feature set. The following are supported if this bit is 1:
      - The compacted format of the extended region of XSAVE areas (see Section 13.4.3).
      - The XSAVEC instruction. If  $EAX[1] = 0$ , execution of XSAVEC causes a #UD.
      - Execution of the compacted form of XRSTOR (see Section 13.8).
    - EAX[2] enumerates support for execution of XGETBV with  $ECX = 1$ . This allows software to determine the state of the init optimization. See Section 13.6.
    - EAX[3] enumerates support for XSAVES, XRSTORS, and the IA32\_XSS MSR. If  $EAX[3] = 0$ , execution of XSAVES or XRSTORS causes a #UD; an attempt to access the IA32\_XSS MSR using RDMSR or WRMSR causes a general-protection exception (#GP). Every processor that supports a supervisor state component sets EAX[3]. Every processor that sets EAX[3] (XSAVES, XRSTORS, IA32\_XSS) will also set EAX[1] (the compaction extensions).
    - EAX[31:4] are reserved.
    - EBX enumerates the size (in bytes) required by the XSAVES instruction for an XSAVE area containing all the state components corresponding to bits currently set in XCR0 | IA32\_XSS.
    - EDX:ECX is a bitmap of all the supervisor state components that can be managed by XSAVES and XRSTORS. A bit can be set in the IA32\_XSS MSR if and only if the corresponding bit is set in this bitmap.

1. If CPUID.1:ECX.XSAVE[bit 26] = 1, XGETBV and XSETBV may be executed with  $ECX = 0$  (to read and write XCR0). Any support for execution of these instructions with other values of ECX is enumerated separately.

**NOTE**

In summary, the XSAVE feature set supports state component  $i$  ( $0 \leq i < 63$ ) if one of the following is true: (1)  $i < 32$  and  $\text{CPUID}.\text{(EAX=0DH,ECX=0):EAX}[i] = 1$ ; (2)  $i \geq 32$  and  $\text{CPUID}.\text{(EAX=0DH,ECX=0):EAX}[i-32] = 1$ ; (3)  $i < 32$  and  $\text{CPUID}.\text{(EAX=0DH,ECX=1):ECX}[i] = 1$ ; or (4)  $i \geq 32$  and  $\text{CPUID}.\text{(EAX=0DH,ECX=1):EDX}[i-32] = 1$ . The XSAVE feature set supports user state component  $i$  if (1) or (2) holds; if (3) or (4) holds, state component  $i$  is a supervisor state component and support is limited to XSAVES and XRSTORS.

- $\text{CPUID}$  function 0DH, sub-function  $i$  ( $i > 1$ ). This sub-function enumerates details for state component  $i$ . If the XSAVE feature set supports state component  $i$  (see note above), the following items provide specific details:
  - EAX enumerates the size (in bytes) required for state component  $i$ .
  - If state component  $i$  is a user state component, EBX enumerates the offset (in bytes, from the base of the XSAVE area) of the section used for state component  $i$ . (This offset applies only when the standard format for the extended region of the XSAVE area is being used; see Section 13.4.3.)
  - If state component  $i$  is a supervisor state component, EBX returns 0.
  - If state component  $i$  is a user state component,  $\text{ECX}[0]$  return 0; if state component  $i$  is a supervisor state component,  $\text{ECX}[0]$  returns 1.
  - The value returned by  $\text{ECX}[1]$  indicates the alignment of state component  $i$  when the compacted format of the extended region of an XSAVE area is used (see Section 13.4.3). If  $\text{ECX}[1]$  returns 0, state component  $i$  is located immediately following the preceding state component; if  $\text{ECX}[1]$  returns 1, state component  $i$  is located on the next 64-byte boundary following the preceding state component.
  - $\text{ECX}[31:2]$  and EDX return 0.

If the XSAVE feature set does not support state component  $i$ , sub-function  $i$  returns 0 in EAX, EBX, ECX, and EDX.

**13.3 ENABLING THE XSAVE FEATURE SET AND XSAVE-ENABLED FEATURES**

Software enables the XSAVE feature set by setting  $\text{CR4.OSXSAVE}[\text{bit } 18]$  to 1 (e.g., with the MOV to CR4 instruction). If this bit is 0, execution of any of XGETBV, XRSTOR, XRSTORS, XSAVE, XSAVEC, XSAVEOPT, XSAVES, and XSETBV causes an invalid-opcode exception (#UD).

When  $\text{CR4.OSXSAVE} = 1$  and  $\text{CPL} = 0$ , executing the XSETBV instruction with  $\text{ECX} = 0$  writes the 64-bit value in  $\text{EDX:EAX}$  to XCR0 (EAX is written to  $\text{XCR0}[31:0]$  and EDX to  $\text{XCR0}[63:32]$ ). (Execution of the XSETBV instruction causes a general-protection fault — #GP — if  $\text{CPL} > 0$ .) The following items provide details regarding individual bits in XCR0:

- $\text{XCR0}[0]$  is associated with x87 state (see Section 13.5.1).  $\text{XCR0}[0]$  is always 1. It has that value coming out of RESET. Executing the XSETBV instruction causes a general-protection fault (#GP) if  $\text{ECX} = 0$  and  $\text{EAX}[0]$  is 0.
- $\text{XCR0}[1]$  is associated with SSE state (see Section 13.5.2). Software can use the XSAVE feature set to manage SSE state only if  $\text{XCR0}[1] = 1$ . The value of  $\text{XCR0}[1]$  in no way determines whether software can execute SSE instructions (these instructions can be executed even if  $\text{XCR0}[1] = 0$ ).

$\text{XCR0}[1]$  is 0 coming out of RESET. As noted in Section 13.2, every processor that supports the XSAVE feature set allows software to set  $\text{XCR0}[1]$ .

- $\text{XCR0}[2]$  is associated with AVX state (see Section 13.5.3). Software can use the XSAVE feature set to manage AVX state only if  $\text{XCR0}[2] = 1$ . In addition, software can execute AVX instructions only if  $\text{CR4.OSXSAVE} = \text{XCR0}[2] = 1$ . Otherwise, any execution of an AVX instruction causes an invalid-opcode exception (#UD).

$\text{XCR0}[2]$  is 0 coming out of RESET. As noted in Section 13.2, a processor allows software to set  $\text{XCR0}[2]$  if and only if  $\text{CPUID}.\text{(EAX=0DH,ECX=0):EAX}[2] = 1$ . In addition, executing the XSETBV instruction causes a general-protection fault (#GP) if  $\text{ECX} = 0$  and  $\text{EAX}[2:1]$  has the value 10b; that is, software cannot enable the XSAVE feature set for AVX state but not for SSE state.

As noted in Section 13.1, the processor will preserve AVX state unmodified if software clears  $\text{XCR0}[2]$ .

However, clearing  $\text{XCR0}[2]$  while AVX state is not in its initial configuration may cause SSE instructions to incur a power and performance penalty. See Section 13.5.3, "Enable the Use Of XSAVE Feature Set And XSAVE State

Components” of *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*, for how system software can avoid this penalty.

- XCR0[4:3] are associated with MPX state (see Section 13.5.4). Software can use the XSAVE feature set to manage MPX state only if XCR0[4:3] = 11b. In addition, software can execute MPX instructions only if CR4.OSXSAVE = 1 and XCR0[4:3] = 11b. Otherwise, any execution of an MPX instruction causes an invalid-opcode exception (#UD).<sup>1</sup>

XCR0[4:3] have value 00b coming out of RESET. As noted in Section 13.2, a processor allows software to set XCR0[4:3] to 11b if and only if CPUID.(EAX=0DH,ECX=0):EAX[4:3] = 11b. In addition, executing the XSETBV instruction causes a general-protection fault (#GP) if ECX = 0, EAX[4:3] is neither 00b nor 11b; that is, software can enable the XSAVE feature set for MPX state only if it does so for both state components.

As noted in Section 13.1, the processor will preserve MPX state unmodified if software clears XCR0[4:3].

- XCR0[7:5] are associated with AVX-512 state (see Section 13.5.5). Software can use the XSAVE feature set to manage AVX-512 state only if XCR0[7:5] = 111b. In addition, software can execute AVX-512 instructions only if CR4.OSXSAVE = 1 and XCR0[7:5] = 111b. Otherwise, any execution of an AVX-512 instruction causes an invalid-opcode exception (#UD).

XCR0[7:5] have value 000b coming out of RESET. As noted in Section 13.2, a processor allows software to set XCR0[7:5] to 111b if and only if CPUID.(EAX=0DH,ECX=0):EAX[7:5] = 111b. In addition, executing the XSETBV instruction causes a general-protection fault (#GP) if ECX = 0, EAX[7:5] is not 000b, and any bit is clear in EAX[2:1] or EAX[7:5]; that is, software can enable the XSAVE feature set for AVX-512 state only if it does so for all three state components, and only if it also does so for AVX state and SSE state. This implies that the value of XCR0[7:5] is always either 000b or 111b.

As noted in Section 13.1, the processor will preserve AVX-512 state unmodified if software clears XCR0[7:5]. However, clearing XCR0[7:5] while AVX-512 state is not in its initial configuration may cause SSE and AVX instructions to incur a power and performance penalty. See Section 13.5.3, “Enable the Use Of XSAVE Feature Set And XSAVE State Components” of *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*, for how system software can avoid this penalty.

- XCR0[9] is associated with PKRU state (see Section 13.5.7). Software can use the XSAVE feature set to manage PKRU state only if XCR0[9] = 1. The value of XCR0[9] in no way determines whether software can use protection keys or execute other instructions that access PKRU state (these instructions can be executed even if XCR0[9] = 0).

XCR0[9] is 0 coming out of RESET. As noted in Section 13.2, a processor allows software to set XCR0[9] if and only if CPUID.(EAX=0DH,ECX=0):EAX[9] = 1.

- XCR0[63:10] and XCR0[8] are reserved.<sup>2</sup> Executing the XSETBV instruction causes a general-protection fault (#GP) if ECX = 0 and any corresponding bit in EDX:EAX is not 0. These bits in XCR0 are all 0 coming out of RESET.

Software operating with CPL > 0 may need to determine whether the XSAVE feature set and certain XSAVE-enabled features have been enabled. If CPL > 0, execution of the MOV from CR4 instruction causes a general-protection fault (#GP). The following alternative mechanisms allow software to discover the enabling of the XSAVE feature set regardless of CPL:

- The value of CR4.OSXSAVE is returned in CPUID.1:ECX.OSXSAVE[bit 27]. If software determines that CPUID.1:ECX.OSXSAVE = 1, the processor supports the XSAVE feature set and the feature set has been enabled in CR4.
- Executing the XGETBV instruction with ECX = 0 returns the value of XCR0 in EDX:EAX. XGETBV can be executed if CR4.OSXSAVE = 1 (if CPUID.1:ECX.OSXSAVE = 1), regardless of CPL.

Thus, software can use the following algorithm to determine the support and enabling for the XSAVE feature set:

1. Use CPUID to discover the value of CPUID.1:ECX.OSXSAVE.
  - If the bit is 0, either the XSAVE feature set is not supported by the processor or has not been enabled by software. Either way, the XSAVE feature set is not available, nor are XSAVE-enabled features such as AVX.

1. If XCR0[3] = 0, executions of CALL, RET, JMP, and Jcc do not initialize the bounds registers.

2. Bit 8 and bit 13 correspond to supervisor state components. Since bits can be set in XCR0 only for user state components, those bits of XCR0 must be 0.

- If the bit is 1, the processor supports the XSAVE feature set — including the XGETBV instruction — and it has been enabled by software. The XSAVE feature set can be used to manage x87 state (because XCR0[0] is always 1). Software requiring more detailed information can go on to the next step.
2. Execute XGETBV with ECX = 0 to discover the value of XCR0. If XCR0[1] = 1, the XSAVE feature set can be used to manage SSE state. If XCR0[2] = 1, the XSAVE feature set can be used to manage AVX state and software can execute AVX instructions. If XCR0[4:3] is 11b, the XSAVE feature set can be used to manage MPX state and software can execute MPX instructions. If XCR0[7:5] is 111b, the XSAVE feature set can be used to manage AVX-512 state and software can execute AVX-512 instructions. If XCR0[9] = 1, the XSAVE feature set can be used to manage PKRU state.

The IA32\_XSS MSR (with MSR index DA0H) is zero coming out of RESET. If CR4.OSXSAVE = 1, CPUID.(EAX=0DH,ECX=1):EAX[3] = 1, and CPL = 0, executing the WRMSR instruction with ECX = DA0H writes the 64-bit value in EDX:EAX to the IA32\_XSS MSR (EAX is written to IA32\_XSS[31:0] and EDX to IA32\_XSS[63:32]). The following items provide details regarding individual bits in the IA32\_XSS MSR:

- IA32\_XSS[8] is associated with PT state (see Section 13.5.6). Software can use XSAVES and XRSTORS to manage PT state only if IA32\_XSS[8] = 1. The value of IA32\_XSS[8] does not determine whether software can use Intel Processor Trace (the feature can be used even if IA32\_XSS[8] = 0).
- IA32\_XSS[13] is associated with HDC state (see Section 13.5.8). Software can use XSAVES and XRSTORS to manage HDC state only if IA32\_XSS[13] = 1. The value of IA32\_XSS[13] does not determine whether software can use hardware duty cycling (the feature can be used even if IA32\_XSS[13] = 0).
- IA32\_XSS[63:14], IA32\_XSS[12:9] and IA32\_XSS[7:0] are reserved.<sup>1</sup> Executing the WRMSR instruction causes a general-protection fault (#GP) if ECX = DA0H and any corresponding bit in EDX:EAX is not 0. These bits in XCR0 are all 0 coming out of RESET.

The IA32\_XSS MSR is 0 coming out of RESET.

There is no mechanism by which software operating with CPL > 0 can discover the value of the IA32\_XSS MSR.

## 13.4 XSAVE AREA

The XSAVE feature set includes instructions that save and restore the XSAVE-managed state components to and from memory: XSAVE, XSAVEOPT, XSAVEC, and XSAVES (for saving); and XRSTOR and XRSTORS (for restoring). The processor organizes the state components in a region of memory called an XSAVE area. Each of the save and restore instructions takes a memory operand that specifies the 64-byte aligned base address of the XSAVE area on which it operates.

Every XSAVE area has the following format:

- The **legacy region**. The legacy region of an XSAVE area comprises the 512 bytes starting at the area's base address. It is used to manage the state components for x87 state and SSE state. The legacy region is described in more detail in Section 13.4.1.
- The **XSAVE header**. The XSAVE header of an XSAVE area comprises the 64 bytes starting at an offset of 512 bytes from the area's base address. The XSAVE header is described in more detail in Section 13.4.2.
- The **extended region**. The extended region of an XSAVE area starts at an offset of 576 bytes from the area's base address. It is used to manage the state components other than those for x87 state and SSE state. The extended region is described in more detail in Section 13.4.3. The size of the extended region is determined by which state components the processor supports and which bits have been set in XCR0 and IA32\_XSS (see Section 13.3).

### 13.4.1 Legacy Region of an XSAVE Area

The legacy region of an XSAVE area comprises the 512 bytes starting at the area's base address. It has the same format as the FXSAVE area (see Section 10.5.1). The XSAVE feature set uses the legacy area for x87 state (state

1. Bit 9 and bits 7:0 correspond to user state components. Since bits can be set in the IA32\_XSS MSR only for supervisor state components, those bits of the MSR must be 0.

component 0) and SSE state (state component 1). Table 13-1 illustrates the format of the first 416 bytes of the legacy region of an XSAVE area.

**Table 13-1. Format of the Legacy Region of an XSAVE Area**

| 15 14                  | 13 12             | 11 10     | 9 8      | 7 6                    | 5                 | 4       | 3 2       | 1 0 |            |
|------------------------|-------------------|-----------|----------|------------------------|-------------------|---------|-----------|-----|------------|
| FIP[63:48] or reserved | FCS or FIP[47:32] | FIP[31:0] |          | FOP                    | Rsvd.             | FTW     | FSW       | FCW | <b>0</b>   |
| MXCSR_MASK             |                   | MXCSR     |          | FDP[63:48] or reserved | FDS or FDP[47:32] |         | FDP[31:0] |     | <b>16</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST0/MM0 |           |     | <b>32</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST1/MM1 |           |     | <b>48</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST2/MM2 |           |     | <b>64</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST3/MM3 |           |     | <b>80</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST4/MM4 |           |     | <b>96</b>  |
| Reserved               |                   |           | Reserved |                        |                   | ST5/MM5 |           |     | <b>112</b> |
| Reserved               |                   |           | Reserved |                        |                   | ST6/MM6 |           |     | <b>128</b> |
| Reserved               |                   |           | Reserved |                        |                   | ST7/MM7 |           |     | <b>144</b> |
| XMM0                   |                   |           |          |                        |                   |         |           |     | <b>160</b> |
| XMM1                   |                   |           |          |                        |                   |         |           |     | <b>176</b> |
| XMM2                   |                   |           |          |                        |                   |         |           |     | <b>192</b> |
| XMM3                   |                   |           |          |                        |                   |         |           |     | <b>208</b> |
| XMM4                   |                   |           |          |                        |                   |         |           |     | <b>224</b> |
| XMM5                   |                   |           |          |                        |                   |         |           |     | <b>240</b> |
| XMM6                   |                   |           |          |                        |                   |         |           |     | <b>256</b> |
| XMM7                   |                   |           |          |                        |                   |         |           |     | <b>272</b> |
| XMM8                   |                   |           |          |                        |                   |         |           |     | <b>288</b> |
| XMM9                   |                   |           |          |                        |                   |         |           |     | <b>304</b> |
| XMM10                  |                   |           |          |                        |                   |         |           |     | <b>320</b> |
| XMM11                  |                   |           |          |                        |                   |         |           |     | <b>336</b> |
| XMM12                  |                   |           |          |                        |                   |         |           |     | <b>352</b> |
| XMM13                  |                   |           |          |                        |                   |         |           |     | <b>368</b> |
| XMM14                  |                   |           |          |                        |                   |         |           |     | <b>384</b> |
| XMM15                  |                   |           |          |                        |                   |         |           |     | <b>400</b> |

The x87 state component comprises bytes 23:0 and bytes 159:32. The SSE state component comprises bytes 31:24 and bytes 415:160. The XSAVE feature set does not use bytes 511:416; bytes 463:416 are reserved.

Section 13.7 through Section 13.9 provide details of how instructions in the XSAVE feature set use the legacy region of an XSAVE area.

### 13.4.2 XSAVE Header

The XSAVE header of an XSAVE area comprises the 64 bytes starting at offset 512 from the area’s base address:

- Bytes 7:0 of the XSAVE header is a state-component bitmap (see Section 13.1) called **XSTATE\_BV**. It identifies the state components in the XSAVE area.

- Bytes 15:8 of the XSAVE header is a state-component bitmap called **XCOMP\_BV**. It is used as follows:
  - **XCOMP\_BV[63]** indicates the format of the extended region of the XSAVE area (see Section 13.4.3). If it is clear, the standard format is used. If it is set, the compacted format is used; **XCOMP\_BV[62:0]** provide format specifics as specified in Section 13.4.3.
  - **XCOMP\_BV[63]** determines which form of the XRSTOR instruction is used. If the bit is set, the compacted form is used; otherwise, the standard form is used. See Section 13.8.
  - All bits in **XCOMP\_BV** should be 0 if the processor does not support the compaction extensions to the XSAVE feature set.
- Bytes 63:16 of the XSAVE header are reserved.

Section 13.7 through Section 13.9 provide details of how instructions in the XSAVE feature set use the XSAVE header of an XSAVE area.

### 13.4.3 Extended Region of an XSAVE Area

The extended region of an XSAVE area starts at byte offset 576 from the area’s base address. The size of the extended region is determined by which state components the processor supports and which bits have been set in **XCRO | IA32\_XSS** (see Section 13.3).

The XSAVE feature set uses the extended area for each state component  $i$ , where  $i \geq 2$ . The following state components are currently supported in the extended area: state component 2 contains AVX state; state components 5–7 contain AVX-512 state; and state component 9 contains PKRU state.

The extended region of the an XSAVE area may have one of two formats. The **standard format** is supported by all processors that support the XSAVE feature set; the **compacted format** is supported by those processors that support the compaction extensions to the XSAVE feature set (see Section 13.2). Bit 63 of the **XCOMP\_BV** field in the XSAVE header (see Section 13.4.2) indicates which format is used.

The following items describe the two possible formats of the extended region:

- **Standard format.** Each state component  $i$  ( $i \geq 2$ ) is located at the byte offset from the base address of the XSAVE area enumerated in **CPUID.(EAX=0DH,ECX=i):EBX**. (**CPUID.(EAX=0DH,ECX=i):EAX** enumerates the number of bytes required for state component  $i$ .)
- **Compacted format.** Each state component  $i$  ( $i \geq 2$ ) is located at a byte offset from the base address of the XSAVE area based on the **XCOMP\_BV** field in the XSAVE header:
  - If **XCOMP\_BV[ $i$ ] = 0**, state component  $i$  is not in the XSAVE area.
  - If **XCOMP\_BV[ $i$ ] = 1**, state component  $i$  is located at a byte offset  $location_i$  from the base address of the XSAVE area, where  $location_i$  is determined by the following items:
    - If **XCOMP\_BV[ $j$ ] = 0** for every  $j$ ,  $2 \leq j < i$ ,  $location_i$  is 576. (This item applies if  $i$  is the first bit set in bits 62:2 of the **XCOMP\_BV**; it implies that state component  $i$  is located at the beginning of the extended region.)
    - Otherwise, let  $j$ ,  $2 \leq j < i$ , be the greatest value such that **XCOMP\_BV[ $j$ ] = 1**. Then  $location_i$  is determined by the following values:  $location_j$ ;  $size_j$ , as enumerated in **CPUID.(EAX=0DH,ECX=j):EAX**; and the value of  $align_i$ , as enumerated in **CPUID.(EAX=0DH,ECX=i):ECX[1]**:
      - If  $align_i = 0$ ,  $location_i = location_j + size_j$ . (This item implies that state component  $i$  is located immediately following the preceding state component whose bit is set in **XCOMP\_BV**.)
      - If  $align_i = 1$ ,  $location_i = \text{ceiling}(location_j + size_j, 64)$ . (This item implies that state component  $i$  is located on the next 64-byte boundary following the preceding state component whose bit is set in **XCOMP\_BV**.)

## 13.5 XSAVE-MANAGED STATE

The section provides details regarding how the XSAVE feature set interacts with the various XSAVE-managed state components.

Unless otherwise stated, the state pertaining to a particular state component is saved beginning at byte 0 of the section of the XSAVE area corresponding to that state component.

### 13.5.1 x87 State

Instructions in the XSAVE feature set can manage the same state of the x87 FPU execution environment (x87 state) that can be managed using the FXSAVE and FXRSTOR instructions. They organize all x87 state as a user state component in the legacy region of the XSAVE area (see Section 13.4.1). This region is illustrated in Table 13-1; the x87 state is listed below, along with details of its interactions with the XSAVE feature set:

- Bytes 1:0, 3:2, 7:6. These are used for the x87 FPU Control Word (FCW), the x87 FPU Status Word (FSW), and the x87 FPU Opcode (FOP), respectively.
- Byte 4 is used for an abridged version of the x87 FPU Tag Word (FTW). The following items describe its usage:
  - For each  $j$ ,  $0 \leq j \leq 7$ , XSAVE, XSAVEOPT, XSAVEC, and XSAVES save a 0 into bit  $j$  of byte 4 if x87 FPU data register  $ST_j$  has an empty tag; otherwise, XSAVE, XSAVEOPT, XSAVEC, and XSAVES save a 1 into bit  $j$  of byte 4.
  - For each  $j$ ,  $0 \leq j \leq 7$ , XRSTOR and XRSTORS establish the tag value for x87 FPU data register  $ST_j$  as follows. If bit  $j$  of byte 4 is 0, the tag for  $ST_j$  in the tag register for that data register is marked empty (11B); otherwise, the x87 FPU sets the tag for  $ST_j$  based on the value being loaded into that register (see below).
- Bytes 15:8 are used as follows:
  - If the instruction has no REX prefix, or if  $REX.W = 0$ :
    - Bytes 11:8 are used for bits 31:0 of the x87 FPU Instruction Pointer Offset (FIP).
    - If  $CPUID.(EAX=07H,ECX=0H):EBX[\text{bit } 13] = 0$ , bytes 13:12 are used for x87 FPU Instruction Pointer Selector (FCS). Otherwise, XSAVE, XSAVEOPT, XSAVEC, and XSAVES save these bytes as 0000H, and XRSTOR and XRSTORS ignore them.
    - Bytes 15:14 are not used.
  - If the instruction has a REX prefix with  $REX.W = 1$ , bytes 15:8 are used for the full 64 bits of FIP.
- Bytes 23:16 are used as follows:
  - If the instruction has no REX prefix, or if  $REX.W = 0$ :
    - Bytes 19:16 are used for bits 31:0 of the x87 FPU Data Pointer Offset (FDP).
    - If  $CPUID.(EAX=07H,ECX=0H):EBX[\text{bit } 13] = 0$ , bytes 21:20 are used for x87 FPU Data Pointer Selector (FDS). Otherwise, XSAVE, XSAVEOPT, XSAVEC, and XSAVES save these bytes as 0000H; and XRSTOR and XRSTORS ignore them.
    - Bytes 23:22 are not used.
  - If the instruction has a REX prefix with  $REX.W = 1$ , bytes 23:16 are used for the full 64 bits of FDP.
- Bytes 31:24 are used for SSE state (see Section 13.5.2).
- Bytes 159:32 are used for the registers  $ST_0$ – $ST_7$  ( $MM_0$ – $MM_7$ ). Each of the 8 registers is allocated a 128-bit region, with the low 80 bits used for the register and the upper 48 bits unused.

x87 state is XSAVE-managed but the x87 FPU feature is not XSAVE-enabled. The XSAVE feature set can operate on x87 state only if the feature set is enabled ( $CR_4.OSXSAVE = 1$ ).<sup>1</sup> Software can otherwise use x87 state even if the XSAVE feature set is not enabled.

### 13.5.2 SSE State

Instructions in the XSAVE feature set can manage the registers used by the streaming SIMD extensions (SSE state) just as the FXSAVE and FXRSTOR instructions do. They organize all SSE state as a user state component in the legacy region of the XSAVE area (see Section 13.4.1). This region is illustrated in Table 13-1; the SSE state is listed below, along with details of its interactions with the XSAVE feature set:

1. The processor ensures that  $XCRO[0]$  is always 1.

- Bytes 23:0 are used for x87 state (see Section 13.5.1).
- Bytes 27:24 are used for the MXCSR register. XRSTOR and XRSTORS generate general-protection faults (#GP) in response to attempts to set any of the reserved bits of the MXCSR register.<sup>1</sup>
- Bytes 31:28 are used for the MXCSR\_MASK value. XRSTOR and XRSTORS ignore this field.
- Bytes 159:32 are used for x87 state.
- Bytes 287:160 are used for the registers XMM0–XMM7.
- Bytes 415:288 are used for the registers XMM8–XMM15. These fields are used only in 64-bit mode. Executions of XSAVE, XSAVEOPT, XSAVEC, and XSAVES outside 64-bit mode do not modify these bytes; executions of XRSTOR and XRSTORS outside 64-bit mode do not update XMM8–XMM15. See Section 13.13.

SSE state is XSAVE-managed but the SSE feature is not XSAVE-enabled. The XSAVE feature set can operate on SSE state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage SSE state (XCR0[1] = 1). Software can otherwise use SSE state even if the XSAVE feature set is not enabled or has not been configured to manage SSE state.

### 13.5.3 AVX State

The register state used by the Intel<sup>®</sup> Advanced Vector Extensions (AVX) comprises the MXCSR register and 16 256-bit vector registers called YMM0–YMM15. The low 128 bits of each register YMM*i* is identical to the SSE register XMM*i*. Thus, the new state register state added by AVX comprises the upper 128 bits of the registers YMM0–YMM15. These 16 128-bit values are denoted YMM0\_H–YMM15\_H and are collectively called **AVX state**.

As noted in Section 13.1, the XSAVE feature set manages AVX state as user state component 2. Thus, AVX state is located in the extended region of the XSAVE area (see Section 13.4.3).

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=2):EBX enumerates the offset (in bytes, from the base of the XSAVE area) of the section of the extended region of the XSAVE area used for AVX state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=2):EAX enumerates the size (in bytes) required for AVX state.

The XSAVE feature set partitions YMM0\_H–YMM15\_H in a manner similar to that used for the XMM registers (see Section 13.5.2). Bytes 127:0 of the AVX-state section are used for YMM0\_H–YMM7\_H. Bytes 255:128 are used for YMM8\_H–YMM15\_H, but they are used only in 64-bit mode. Executions of XSAVE, XSAVEOPT, XSAVEC, and XSAVES outside 64-bit mode do not modify bytes 255:128; executions of XRSTOR and XRSTORS outside 64-bit mode do not update YMM8\_H–YMM15\_H. See Section 13.13. In general, bytes 16/+15:16*i* are used for YMM*i*\_H (for 0 ≤ *i* ≤ 15).

AVX state is XSAVE-managed and the AVX feature is XSAVE-enabled. The XSAVE feature set can operate on AVX state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage AVX state (XCR0[2] = 1). AVX instructions cannot be used unless the XSAVE feature set is enabled and has been configured to manage AVX state.

### 13.5.4 MPX State

The register state used by the Intel<sup>®</sup> Memory Protection Extensions (MPX) comprises the 4 128-bit bounds registers BND0–BND3 (**BNDREGS state**); and the 64-bit user-mode configuration register BNDCFGU and the 64-bit MPX status register BNDSTATUS (collectively, **BNDCSR state**). Together, these two user state components compose **MPX state**.

As noted in Section 13.1, the XSAVE feature set manages MPX state as state components 3–4. Thus, MPX state is located in the extended region of the XSAVE area (see Section 13.4.3). The following items detail how these state components are organized in this region:

- **BNDREGS state.**  
As noted in Section 13.2, CPUID.(EAX=0DH,ECX=3):EBX enumerates the offset (in bytes, from the base of the XSAVE area) of the section of the extended region of the XSAVE area used for BNDREGS state (when the

1. While MXCSR and MXCSR\_MASK are part of SSE state, their treatment by the XSAVE feature set is not the same as that of the XMM registers. See Section 13.7 through Section 13.11 for details.



standard format of the extended region is used). CPUID.(EAX=0DH,ECX=3):EAX enumerates the size (in bytes) required for BNDREGS state. The BNDREGS section is used for the 4 128-bit bound registers BND0–BND3, with bytes  $16i+15:16i$  being used for BND*i*.

- **BNDCSR state.**

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=4):EBX enumerates the offset of the section of the extended region of the XSAVE area used for BNDCSR state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=4):EAX enumerates the size (in bytes) required for BNDCSR state. In the BNDSCR section, bytes 7:0 are used for BNDCFGU and bytes 15:8 are used for BNDSTATUS.

Both components of MPX state are XSAVE-managed and the MPX feature is XSAVE-enabled. The XSAVE feature set can operate on MPX state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage MPX state (XCR0[4:3] = 11b). MPX instructions cannot be used unless the XSAVE feature set is enabled and has been configured to manage MPX state.

### 13.5.5 AVX-512 State

The register state used by the Intel® Advanced Vector Extensions 512 (AVX-512) comprises the MXCSR register, the 8 64-bit opmask registers k0–k7, and 32 512-bit vector registers called ZMM0–ZMM31. For each *i*,  $0 \leq i \leq 15$ , the low 256 bits of register ZMM*i* is identical to the AVX register YMM*i*. Thus, the new state register state added by AVX comprises the following user state components:

- The opmask registers, collectively called **opmask state**.
- The upper 256 bits of the registers ZMM0–ZMM15. These 16 256-bit values are denoted ZMM0\_H–ZMM15\_H and are collectively called **ZMM\_Hi256 state**.
- The 16 512-bit registers ZMM16–ZMM31, collectively called **Hi16\_ZMM state**.

Together, these three state components compose **AVX-512 state**.

As noted in Section 13.1, the XSAVE feature set manages AVX-512 state as state components 5–7. Thus, AVX-512 state is located in the extended region of the XSAVE area (see Section 13.4.3). The following items detail how these state components are organized in this region:

- **Opmask state.**

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=5):EBX enumerates the offset (in bytes, from the base of the XSAVE area) of the section of the extended region of the XSAVE area used for opmask state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=5):EAX enumerates the size (in bytes) required for opmask state. The opmask section is used for the 8 64-bit bound registers k0–k7, with bytes  $8i+7:8i$  being used for k*i*.

- **ZMM\_Hi256 state.**

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=6):EBX enumerates the offset of the section of the extended region of the XSAVE area used for ZMM\_Hi256 state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=6):EAX enumerates the size (in bytes) required for ZMM\_Hi256 state.

The XSAVE feature set partitions ZMM0\_H–ZMM15\_H in a manner similar to that used for the XMM registers (see Section 13.5.2). Bytes 255:0 of the ZMM\_Hi256-state section are used for ZMM0\_H–ZMM7\_H. Bytes 511:256 are used for ZMM8\_H–ZMM15\_H, but they are used only in 64-bit mode. Executions of XSAVE, XSAVEOPT, XSAVEC, and XSAVES outside 64-bit mode do not modify bytes 511:256; executions of XRSTOR and XRSTORS outside 64-bit mode do not update ZMM8\_H–ZMM15\_H. See Section 13.13. In general, bytes  $32i+31:32i$  are used for ZMM*i*\_H (for  $0 \leq i \leq 15$ ).

- **Hi16\_ZMM state.**

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=7):EBX enumerates the offset of the section of the extended region of the XSAVE area used for Hi16\_ZMM state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=7):EAX enumerates the size (in bytes) required for Hi16\_ZMM state.

The XSAVE feature set accesses Hi16\_ZMM state only in 64-bit mode. Executions of XSAVE, XSAVEOPT, XSAVEC, and XSAVES outside 64-bit mode do not modify the Hi16\_ZMM section; executions of XRSTOR and XRSTORS outside 64-bit mode do not update ZMM16–ZMM31. See Section 13.13. In general, bytes  $64(i-16)+63:64(i-16)$  are used for ZMM*i* (for  $16 \leq i \leq 31$ ).

All three components of AVX-512 state are XSAVE-managed and the AVX-512 feature is XSAVE-enabled. The XSAVE feature set can operate on AVX-512 state only if the feature set is enabled (CR4.OSXSAVE = 1) and has

been configured to manage AVX-512 state (XCR0[7:5] = 111b). AVX-512 instructions cannot be used unless the XSAVE feature set is enabled and has been configured to manage AVX-512 state.

### 13.5.6 PT State

The register state used by Intel Processor Trace (PT state) comprises the following 9 MSRs: IA32\_RTIT\_CTL, IA32\_RTIT\_OUTPUT\_BASE, IA32\_RTIT\_OUTPUT\_MASK\_PTRS, IA32\_RTIT\_STATUS, IA32\_RTIT\_CR3\_MATCH, IA32\_RTIT\_ADDR0\_A, IA32\_RTIT\_ADDR0\_B, IA32\_RTIT\_ADDR1\_A, and IA32\_RTIT\_ADDR1\_B.<sup>1</sup>

As noted in Section 13.1, the XSAVE feature set manages PT state as supervisor state component 8. Thus, PT state is located in the extended region of the XSAVE area (see Section 13.4.3). As noted in Section 13.2, CPUID.(EAX=0DH,ECX=8):EAX enumerates the size (in bytes) required for PT state. The MSRs are each allocated 8 bytes in the state component in the order given above. Thus, IA32\_RTIT\_CTL is at byte offset 0, IA32\_RTIT\_OUTPUT\_BASE at byte offset 8, etc. Any locations in the state component at or beyond byte offset 72 are reserved.

PT state is XSAVE-managed but Intel Processor Trace is not XSAVE-enabled. The XSAVE feature set can operate on PT state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage PT state (IA32\_XSS[8] = 1). Software can otherwise use Intel Processor Trace and access its MSRs (using RDMSR and WRMSR) even if the XSAVE feature set is not enabled or has not been configured to manage PT state.

The following items describe special treatment of PT state by the XSAVES and XRSTORS instructions:

- If XSAVES saves PT state, the instruction clears IA32\_RTIT\_CTL.TraceEn (bit 0) after saving the value of the IA32\_RTIT\_CTL MSR and before saving any other PT state. If XSAVES causes a fault or a VM exit, it restores IA32\_RTIT\_CTL.TraceEn to its original value.
- If XSAVES saves PT state, the instruction saves zeroes in the reserved portions of the state component.
- If XRSTORS would restore (or initialize) PT state and IA32\_RTIT\_CTL.TraceEn = 1, the instruction causes a general-protection exception (#GP) before modifying PT state.
- If XRSTORS causes an exception or a VM exit, it does so before any modification to IA32\_RTIT\_CTL.TraceEn (even if it has loaded other PT state).

### 13.5.7 PKRU State

The register state used by the protection-key feature (PKRU state) is the 32-bit PKRU register. As noted in Section 13.1, the XSAVE feature set manages PKRU state as user state component 9. Thus, PKRU state is located in the extended region of the XSAVE area (see Section 13.4.3).

As noted in Section 13.2, CPUID.(EAX=0DH,ECX=9):EBX enumerates the offset (in bytes, from the base of the XSAVE area) of the section of the extended region of the XSAVE area used for PKRU state (when the standard format of the extended region is used). CPUID.(EAX=0DH,ECX=9):EAX enumerates the size (in bytes) required for PKRU state. The XSAVE feature set uses bytes 3:0 of the PK-state section for the PKRU register.

PKRU state is XSAVE-managed but the protection-key feature is not XSAVE-enabled. The XSAVE feature set can operate on PKRU state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage PKRU state (XCR0[9] = 1). Software can otherwise use protection keys and access PKRU state even if the XSAVE feature set is not enabled or has not been configured to manage PKRU state.

The value of the PKRU register determines the access rights for user-mode linear addresses. (See Section 4.6, "Access Rights," of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.) The access rights that pertain to an execution of the XRSTOR and XRSTORS instructions are determined by the value of the register before the execution and not by any value that the execution might load into the PKRU register.

---

1. These MSRs might not be supported by every processor that supports Intel Processor Trace. Software can use the CPUID instruction to discover which are supported; see Section 35.3.1, "Detection of Intel Processor Trace and Capability Enumeration," of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C*.

### 13.5.8 HDC State

The register state used by hardware duty cycling (HDC state) comprises the IA32\_PM\_CTL1 MSR.

As noted in Section 13.1, the XSAVE feature set manages HDC state as supervisor state component 13. Thus, HDC state is located in the extended region of the XSAVE area (see Section 13.4.3). As noted in Section 13.2, CPUID.(EAX=0DH,ECX=13):EAX enumerates the size (in bytes) required for PT state. The IA32\_PM\_CTL1 MSR is allocated 8 bytes at byte offset 0 in the state component.

HDC state is XSAVE-managed but hardware duty cycling is not XSAVE-enabled. The XSAVE feature set can operate on HDC state only if the feature set is enabled (CR4.OSXSAVE = 1) and has been configured to manage HDC state (IA32\_XSS[13] = 1). Software can otherwise use hardware duty cycle and access the IA32\_PM\_CTL1 MSR (using RDMSR and WRMSR) even if the XSAVE feature set is not enabled or has not been configured to manage HDC state.

## 13.6 PROCESSOR TRACKING OF XSAVE-MANAGED STATE

The XSAVEOPT, XSAVEC, and XSAVES instructions use two optimizations to reduce the amount of data that they write to memory. They avoid writing data for any state component known to be in its initial configuration (the **init optimization**). In addition, if either XSAVEOPT or XSAVES is using the same XSAVE area as that used by the most recent execution of XRSTOR or XRSTORS, it may avoid writing data for any state component whose configuration is known not to have been modified since then (the **modified optimization**). (XSAVE does not use these optimizations, and XSAVEC does not use the modified optimization.) The operation of XSAVEOPT, XSAVEC, and XSAVES are described in more detail in Section 13.9 through Section 13.11.

A processor can support the init and modified optimizations with special hardware that tracks the state components that might benefit from those optimizations. Other implementations might not include such hardware; such a processor would always consider each such state component as not in its initial configuration and as modified since the last execution of XRSTOR or XRSTORS.

The following notation describes the state of the init and modified optimizations:

- XINUSE denotes the state-component bitmap corresponding to the init optimization. If  $XINUSE[i] = 0$ , state component  $i$  is known to be in its initial configuration; otherwise  $XINUSE[i] = 1$ . It is possible for  $XINUSE[i]$  to be 1 even when state component  $i$  is in its initial configuration. On a processor that does not support the init optimization,  $XINUSE[i]$  is always 1 for every value of  $i$ .

Executing XGETBV with ECX = 1 returns in EDX:EAX the logical-AND of XCR0 and the current value of the XINUSE state-component bitmap. Such an execution of XGETBV always sets EAX[1] to 1 if XCR0[1] = 1 and MXCSR does not have its RESET value of 1F80H. Section 13.2 explains how software can determine whether a processor supports this use of XGETBV.

- XMODIFIED denotes the state-component bitmap corresponding to the modified optimization. If  $XMODIFIED[i] = 0$ , state component  $i$  is known not to have been modified since the most recent execution of XRSTOR or XRSTORS; otherwise  $XMODIFIED[i] = 1$ . It is possible for  $XMODIFIED[i]$  to be 1 even when state component  $i$  has not been modified since the most recent execution of XRSTOR or XRSTORS. On a processor that does not support the modified optimization,  $XMODIFIED[i]$  is always 1 for every value of  $i$ .

A processor that implements the modified optimization saves information about the most recent execution of XRSTOR or XRSTORS in a quantity called XRSTOR\_INFO, a 4-tuple containing the following: (1) the CPL; (2) whether the logical processor was in VMX non-root operation; (3) the linear address of the XSAVE area; and (4) the XCOMP\_BV field in the XSAVE area. An execution of XSAVEOPT or XSAVES uses the modified optimization only if that execution corresponds to XRSTOR\_INFO on these four parameters.

This mechanism implies that, depending on details of the operating system, the processor might determine that an execution of XSAVEOPT by one user application corresponds to an earlier execution of XRSTOR by a different application. For this reason, Intel recommends the application software not use the XSAVEOPT instruction.

The following items specify the initial configuration each state component (for the purposes of defining the XINUSE bitmap):

- **x87 state.** x87 state is in its initial configuration if the following all hold: FCW is 037FH; FSW is 0000H; FTW is FFFFH; FCS and FDS are each 0000H; FIP and FDP are each 00000000\_00000000H; each of ST0–ST7 is 0000\_00000000\_00000000H.

- **SSE state.** In 64-bit mode, SSE state is in its initial configuration if each of XMM0–XMM15 is 0. Outside 64-bit mode, SSE state is in its initial configuration if each of XMM0–XMM7 is 0. XINUSE[1] pertains only to the state of the XMM registers and not to MXCSR. An execution of XRSTOR or XRSTORS outside 64-bit mode does not update XMM8–XMM15. (See Section 13.13.)
- **AVX state.** In 64-bit mode, AVX state is in its initial configuration if each of YMM0\_H–YMM15\_H is 0. Outside 64-bit mode, AVX state is in its initial configuration if each of YMM0\_H–YMM7\_H is 0. An execution of XRSTOR or XRSTORS outside 64-bit mode does not update YMM8\_H–YMM15\_H. (See Section 13.13.)
- **BNDREGS state.** BNDREGS state is in its initial configuration if the value of each of BND0–BND3 is 0.
- **BNDCSR state.** BNDCSR state is in its initial configuration if BNDCFGU and BNDCSR each has value 0.
- **Opmask state.** Opmask state is in its initial configuration if each of the opmask registers k0–k7 is 0.
- **ZMM\_Hi256 state.** In 64-bit mode, ZMM\_Hi256 state is in its initial configuration if each of ZMM0\_H–ZMM15\_H is 0. Outside 64-bit mode, ZMM\_Hi256 state is in its initial configuration if each of ZMM0\_H–ZMM7\_H is 0. An execution of XRSTOR or XRSTORS outside 64-bit mode does not update ZMM8\_H–ZMM15\_H. (See Section 13.13.)
- **Hi16\_ZMM state.** In 64-bit mode, Hi16\_ZMM state is in its initial configuration if each of ZMM16–ZMM31 is 0. Outside 64-bit mode, Hi16\_ZMM state is always in its initial configuration. An execution of XRSTOR or XRSTORS outside 64-bit mode does not update ZMM31–ZMM31. (See Section 13.13.)
- **PT state.** PT state is in its initial configuration if each of the 9 MSRs is 0.
- **PKRU state.** PKRU state is in its initial configuration if the value of the PKRU is 0.
- **HDC state.** HDC state is in its initial configuration if the value of the IA32\_PM\_CTL1 MSR is 0.

## 13.7 OPERATION OF XSAVE

The XSAVE instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair EDX:EAX is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**. The logical-AND of XCR0 and the instruction mask is the **requested-feature bitmap (RFBM)** of the user state components to be saved.

The following conditions cause execution of the XSAVE instruction to generate a fault:

- If the XSAVE feature set is not enabled ( $CR4.OSXSAVE = 0$ ), an invalid-opcode exception (#UD) occurs.
- If  $CR0.TS[\text{bit } 3]$  is 1, a device-not-available exception (#NM) occurs.
- If the address of the XSAVE area is not 64-byte aligned, a general-protection exception (#GP) occurs.<sup>1</sup>

If none of these conditions cause a fault, execution of XSAVE reads the XSTATE\_BV field of the XSAVE header (see Section 13.4.2) and writes it back to memory, setting  $XSTATE\_BV[i]$  ( $0 \leq i \leq 63$ ) as follows:

- If  $RFBM[i] = 0$ ,  $XSTATE\_BV[i]$  is not changed.
- If  $RFBM[i] = 1$ ,  $XSTATE\_BV[i]$  is set to the value of  $XINUSE[i]$ . Section 13.6 defines XINUSE to describe the processor init optimization and specifies the initial configuration of each state component. The nature of that optimization implies the following:
  - If state component  $i$  is in its initial configuration,  $XINUSE[i]$  may be either 0 or 1, and  $XSTATE\_BV[i]$  may be written with either 0 or 1.  
 XINUSE[1] pertains only to the state of the XMM registers and not to MXCSR. Thus,  $XSTATE\_BV[1]$  may be written with 0 even if MXCSR does not have its RESET value of 1F80H.
  - If state component  $i$  is not in its initial configuration,  $XINUSE[i] = 1$  and  $XSTATE\_BV[i]$  is written with 1.  
 (As explained in Section 13.6, the initial configurations of some state components may depend on whether the processor is in 64-bit mode.)

The XSAVE instruction does not write any part of the XSAVE header other than the XSTATE\_BV field; in particular, it does not write to the XCOMP\_BV field.

1. If  $CR0.AM = 1$ ,  $CPL = 3$ , and  $EFLAGS.AC = 1$ , an alignment-check exception (#AC) may occur instead of #GP.

Execution of XSAVE saves into the XSAVE area those state components corresponding to bits that are set in RFBM. State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the XSAVE instruction always uses the standard format for the extended region (see Section 13.4.3).

The MXCSR register and MXCSR\_MASK are part of SSE state (see Section 13.5.2) and are thus associated with RFBM[1]. However, the XSAVE instruction also saves these values when RFBM[2] = 1 (even if RFBM[1] = 0).

See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes. See Section 13.13 for details regarding faults caused by memory accesses.

## 13.8 OPERATION OF XRSTOR

The XRSTOR instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair EDX:EAX is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**. The logical-AND of XCR0 and the instruction mask is the **requested-feature bitmap (RFBM)** of the user state components to be restored.

The following conditions cause execution of the XRSTOR instruction to generate a fault:

- If the XSAVE feature set is not enabled (CR4.OSXSAVE = 0), an invalid-opcode exception (#UD) occurs.
- If CR0.TS[bit 3] is 1, a device-not-available exception (#NM) occurs.
- If the address of the XSAVE area is not 64-byte aligned, a general-protection exception (#GP) occurs.<sup>1</sup>

After checking for these faults, the XRSTOR instruction reads the XCOMP\_BV field in the XSAVE area's XSAVE header (see Section 13.4.2). If XCOMP\_BV[63] = 0, the **standard form of XRSTOR** is executed (see Section 13.8.1); otherwise, the **compacted form of XRSTOR** is executed (see Section 13.8.2).<sup>2</sup>

See Section 13.2 for details of how to determine whether the compacted form of XRSTOR is supported.

### 13.8.1 Standard Form of XRSTOR

The standard form of XRSTOR performs additional fault checking. Either of the following conditions causes a general-protection exception (#GP):

- The XSTATE\_BV field of the XSAVE header sets a bit that is not set in XCR0.
- Bytes 23:8 of the XSAVE header are not all 0 (this implies that all bits in XCOMP\_BV are 0).<sup>3</sup>

If none of these conditions cause a fault, the processor updates each state component  $i$  for which RFBM[ $i$ ] = 1. XRSTOR updates state component  $i$  based on the value of bit  $i$  in the XSTATE\_BV field of the XSAVE header:

- If XSTATE\_BV[ $i$ ] = 0, the state component is set to its initial configuration. Section 13.6 specifies the initial configuration of each state component.

The initial configuration of state component 1 pertains only to the XMM registers and not to MXCSR. See below for the treatment of MXCSR

- If XSTATE\_BV[ $i$ ] = 1, the state component is loaded with data from the XSAVE area. See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes. See Section 13.13 for details regarding faults caused by memory accesses.

1. If CR0.AM = 1, CPL = 3, and EFLAGS.AC = 1, an alignment-check exception (#AC) may occur instead of #GP.

2. If the processor does not support the compacted form of XRSTOR, it may execute the standard form of XRSTOR without first reading the XCOMP\_BV field. A processor supports the compacted form of XRSTOR only if it enumerates CPUID.(EAX=0DH,ECX=1):EAX[1] as 1.

3. Bytes 63:24 of the XSAVE header are also reserved. Software should ensure that bytes 63:16 of the XSAVE header are all 0 in any XSAVE area. (Bytes 15:8 should also be 0 if the XSAVE area is to be used on a processor that does not support the compaction extensions to the XSAVE feature set.)

State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the standard form of XRSTOR uses the standard format for the extended region (see Section 13.4.3).

The MXCSR register is part of state component 1, SSE state (see Section 13.5.2). However, the standard form of XRSTOR loads the MXCSR register from memory whenever the RFBM[1] (SSE) or RFBM[2] (AVX) is set, regardless of the values of XSTATE\_BV[1] and XSTATE\_BV[2]. The standard form of XRSTOR causes a general-protection exception (#GP) if it would load MXCSR with an illegal value.

### 13.8.2 Compacted Form of XRSTOR

The compacted form of XRSTOR performs additional fault checking. Any of the following conditions causes a #GP:

- The XCOMP\_BV field of the XSAVE header sets a bit in the range 62:0 that is not set in XCR0.
- The XSTATE\_BV field of the XSAVE header sets a bit (including bit 63) that is not set in XCOMP\_BV.
- Bytes 63:16 of the XSAVE header are not all 0.

If none of these conditions cause a fault, the processor updates each state component  $i$  for which RFBM[ $i$ ] = 1. XRSTOR updates state component  $i$  based on the value of bit  $i$  in the XSTATE\_BV field of the XSAVE header:

- If XSTATE\_BV[ $i$ ] = 0, the state component is set to its initial configuration. Section 13.6 specifies the initial configuration of each state component.

If XSTATE\_BV[1] = 0, the compacted form XRSTOR initializes MXCSR to 1F80H. (This differs from the standard form of XRSTOR, which loads MXCSR from the XSAVE area whenever either RFBM[1] or RFBM[2] is set.)

State component  $i$  is set to its initial configuration as indicated above if RFBM[ $i$ ] = 1 and XSTATE\_BV[ $i$ ] = 0 — even if XCOMP\_BV[ $i$ ] = 0. This is true for all values of  $i$ , including 0 (x87 state) and 1 (SSE state).

- If XSTATE\_BV[ $i$ ] = 1, the state component is loaded with data from the XSAVE area.<sup>1</sup> See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes. See Section 13.13 for details regarding faults caused by memory accesses.

State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the compacted form of the XRSTOR instruction uses the compacted format for the extended region (see Section 13.4.3).

The MXCSR register is part of SSE state (see Section 13.5.2) and is thus loaded from memory if RFBM[1] = XSTATE\_BV[1] = 1. The compacted form of XRSTOR does not consider RFBM[2] (AVX) when determining whether to update MXCSR. (This is a difference from the standard form of XRSTOR.) The compacted form of XRSTOR causes a general-protection exception (#GP) if it would load MXCSR with an illegal value.

### 13.8.3 XRSTOR and the Init and Modified Optimizations

Execution of the XRSTOR instruction causes the processor to update its tracking for the init and modified optimizations (see Section 13.6). The following items provide details:

- The processor updates its tracking for the init optimization as follows:
  - If RFBM[ $i$ ] = 0, XINUSE[ $i$ ] is not changed.
  - If RFBM[ $i$ ] = 1 and XSTATE\_BV[ $i$ ] = 0, state component  $i$  may be tracked as init; XINUSE[ $i$ ] may be set to 0 or 1. (As noted in Section 13.6, a processor need not implement the init optimization for state component  $i$ ; a processor that does not do so implicitly maintains XINUSE[ $i$ ] = 1 at all times.)
  - If RFBM[ $i$ ] = 1 and XSTATE\_BV[ $i$ ] = 1, state component  $i$  is tracked as not init; XINUSE[ $i$ ] is set to 1.
- The processor updates its tracking for the modified optimization and records information about the XRSTOR execution for future interaction with the XSAVEOPT and XSAVES instructions (see Section 13.9 and Section 13.11) as follows:
  - If RFBM[ $i$ ] = 0, state component  $i$  is tracked as modified; XMODIFIED[ $i$ ] is set to 1.

1. Earlier fault checking ensured that, if the instruction has reached this point in execution and XSTATE\_BV[ $i$ ] is 1, then XCOMP\_BV[ $i$ ] is also 1.

- If  $RFBM[i] = 1$ , state component  $i$  may be tracked as unmodified;  $XMODIFIED[i]$  may be set to 0 or 1. (As noted in Section 13.6, a processor need not implement the modified optimization for state component  $i$ ; a processor that does not do so implicitly maintains  $XMODIFIED[i] = 1$  at all times.)
- $XRSTOR\_INFO$  is set to the 4-tuple  $\langle w, x, y, z \rangle$ , where  $w$  is the CPL (0);  $x$  is 1 if the logical processor is in VMX non-root operation and 0 otherwise;  $y$  is the linear address of the XSAVE area; and  $z$  is  $XCOMP\_BV$ . In particular, the standard form of  $XRSTOR$  always sets  $z$  to all zeroes, while the compacted form of  $XRSTORS$  never does so (because it sets at least bit 63 to 1).

## 13.9 OPERATION OF XSAVEOPT

The operation of  $XSAVEOPT$  is similar to that of  $XSAVE$ . Unlike  $XSAVE$ ,  $XSAVEOPT$  uses the init optimization (by which it may omit saving state components that are in their initial configuration) and the modified optimization (by which it may omit saving state components that have not been modified since the last execution of  $XRSTOR$ ); see Section 13.6. See Section 13.2 for details of how to determine whether  $XSAVEOPT$  is supported.

The  $XSAVEOPT$  instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair  $EDX:EAX$  is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**. The logical (bitwise) AND of  $XCR0$  and the instruction mask is the **requested-feature bitmap (RFBM)** of the user state components to be saved.

The following conditions cause execution of the  $XSAVEOPT$  instruction to generate a fault:

- If the XSAVE feature set is not enabled ( $CR4.OSXSAVE = 0$ ), an invalid-opcode exception ( $\#UD$ ) occurs.
- If  $CR0.TS[\text{bit } 3]$  is 1, a device-not-available exception ( $\#NM$ ) occurs.
- If the address of the XSAVE area is not 64-byte aligned, a general-protection exception ( $\#GP$ ) occurs.<sup>1</sup>

If none of these conditions cause a fault, execution of  $XSAVEOPT$  reads the  $XSTATE\_BV$  field of the XSAVE header (see Section 13.4.2) and writes it back to memory, setting  $XSTATE\_BV[i]$  ( $0 \leq i \leq 63$ ) as follows:

- If  $RFBM[i] = 0$ ,  $XSTATE\_BV[i]$  is not changed.
- If  $RFBM[i] = 1$ ,  $XSTATE\_BV[i]$  is set to the value of  $XINUSE[i]$ . Section 13.6 defines  $XINUSE$  to describe the processor init optimization and specifies the initial configuration of each state component. The nature of that optimization implies the following:
  - If the state component is in its initial configuration,  $XINUSE[i]$  may be either 0 or 1, and  $XSTATE\_BV[i]$  may be written with either 0 or 1.
 

$XINUSE[1]$  pertains only to the state of the XMM registers and not to  $MXCSR$ . Thus,  $XSTATE\_BV[1]$  may be written with 0 even if  $MXCSR$  does not have its  $RESET$  value of  $1F80H$ .
  - If the state component is not in its initial configuration,  $XSTATE\_BV[i]$  is written with 1.

(As explained in Section 13.6, the initial configurations of some state components may depend on whether the processor is in 64-bit mode.)

The  $XSAVEOPT$  instruction does not write any part of the XSAVE header other than the  $XSTATE\_BV$  field; in particular, it does not write to the  $XCOMP\_BV$  field.

Execution of  $XSAVEOPT$  saves into the XSAVE area those state components corresponding to bits that are set in RFBM (subject to the optimizations described below). State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the  $XSAVEOPT$  instruction always uses the standard format for the extended region (see Section 13.4.3).

See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes. See Section 13.13 for details regarding faults caused by memory accesses.

Execution of  $XSAVEOPT$  performs two optimizations that reduce the amount of data written to memory:

---

1. If  $CR0.AM = 1$ ,  $CPL = 3$ , and  $EFLAGS.AC = 1$ , an alignment-check exception ( $\#AC$ ) may occur instead of  $\#GP$ .

- **Init optimization.**

If  $XINUSE[i] = 0$ , state component  $i$  is not saved to the XSAVE area (even if  $RFBM[i] = 1$ ). (See below for exceptions made for MXCSR.)

- **Modified optimization.**

Each execution of XRSTOR and XRSTORS establishes XRSTOR\_INFO as a 4-tuple  $\langle w, x, y, z \rangle$  (see Section 13.8.3 and Section 13.12). Execution of XSAVEOPT uses the modified optimization only if the following all hold for the current value of XRSTOR\_INFO:

- $w = \text{CPL}$ ;
- $x = 1$  if and only if the logical processor is in VMX non-root operation;
- $y$  is the linear address of the XSAVE area being used by XSAVEOPT; and
- $z$  is 00000000\_00000000H. (This last item implies that XSAVEOPT does not use the modified optimization if the last execution of XRSTOR used the compacted form, or if an execution of XRSTORS followed the last execution of XRSTOR.)

If XSAVEOPT uses the modified optimization and  $XMODIFIED[i] = 0$  (see Section 13.6), state component  $i$  is not saved to the XSAVE area.

(In practice, the benefit of the modified optimization for state component  $i$  depends on how the processor is tracking state component  $i$ ; see Section 13.6. Limitations on the tracking ability may result in state component  $i$  being saved even though it is in the same configuration that was loaded by the previous execution of XRSTOR.)

Depending on details of the operating system, an execution of XSAVEOPT by a user application might use the modified optimization when the most recent execution of XRSTOR was by a different application. Because of this, Intel recommends the application software not use the XSAVEOPT instruction.

The MXCSR register and MXCSR\_MASK are part of SSE state (see Section 13.5.2) and are thus associated with bit 1 of RFBM. However, the XSAVEOPT instruction also saves these values when  $RFBM[2] = 1$  (even if  $RFBM[1] = 0$ ). The init and modified optimizations do not apply to the MXCSR register and MXCSR\_MASK.

## 13.10 OPERATION OF XSAVEC

The operation of XSAVEC is similar to that of XSAVE. Two main differences are (1) XSAVEC uses the compacted format for the extended region of the XSAVE area; and (2) XSAVEC uses the init optimization (see Section 13.6). Unlike XSAVEOPT, XSAVEC does not use the modified optimization. See Section 13.2 for details of how to determine whether XSAVEC is supported.

The XSAVEC instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair EDX:EAX is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**. The logical (bitwise) AND of XCR0 and the instruction mask is the **requested-feature bitmap (RFBM)** of the user state components to be saved.

The following conditions cause execution of the XSAVEC instruction to generate a fault:

- If the XSAVE feature set is not enabled ( $CR4.OSXSAVE = 0$ ), an invalid-opcode exception (#UD) occurs.
- If  $CR0.TS[\text{bit } 3]$  is 1, a device-not-available exception (#NM) occurs.
- If the address of the XSAVE area is not 64-byte aligned, a general-protection exception (#GP) occurs.<sup>1</sup>

If none of these conditions cause a fault, execution of XSAVEC writes the XSTATE\_BV field of the XSAVE header (see Section 13.4.2), setting  $XSTATE\_BV[i]$  ( $0 \leq i \leq 63$ ) as follows:<sup>2</sup>

- If  $RFBM[i] = 0$ ,  $XSTATE\_BV[i]$  is written as 0.
- If  $RFBM[i] = 1$ ,  $XSTATE\_BV[i]$  is set to the value of  $XINUSE[i]$  (see below for an exception made for  $XSTATE\_BV[1]$ ). Section 13.6 defines XINUSE to describe the processor init optimization and specifies the initial configuration of each state component. The nature of that optimization implies the following:
  - If state component  $i$  is in its initial configuration,  $XSTATE\_BV[i]$  may be written with either 0 or 1.

1. If  $CR0.AM = 1$ ,  $CPL = 3$ , and  $EFLAGS.AC = 1$ , an alignment-check exception (#AC) may occur instead of #GP.

2. Unlike the XSAVE and XSAVEOPT instructions, the XSAVEC instruction does **not** read the XSTATE\_BV field of the XSAVE header.



- If state component  $i$  is not in its initial configuration, `XSTATE_BV[ $i$ ]` is written with 1.

`XINUSE[1]` pertains only to the state of the XMM registers and not to MXCSR. However, if `RFBM[1] = 1` and MXCSR does not have the value 1F80H, XSAVEC writes `XSTATE_BV[1]` as 1 even if `XINUSE[1] = 0`.

(As explained in Section 13.6, the initial configurations of some state components may depend on whether the processor is in 64-bit mode.)

The XSAVEC instruction sets bit 63 of the `XCOMP_BV` field of the XSAVE header while writing `RFBM[62:0]` to `XCOMP_BV[62:0]`. The XSAVEC instruction does not write any part of the XSAVE header other than the `XSTATE_BV` and `XCOMP_BV` fields.

Execution of XSAVEC saves into the XSAVE area those state components corresponding to bits that are set in `RFBM` (subject to the init optimization described below). State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the XSAVEC instruction always uses the compacted format for the extended region (see Section 13.4.3).

See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes. See Section 13.13 for details regarding faults caused by memory accesses.

Execution of XSAVEC performs the init optimization to reduce the amount of data written to memory. If `XINUSE[ $i$ ] = 0`, state component  $i$  is not saved to the XSAVE area (even if `RFBM[ $i$ ] = 1`). However, if `RFBM[1] = 1` and MXCSR does not have the value 1F80H, XSAVEC writes saves all of state component 1 (SSE — including the XMM registers) even if `XINUSE[1] = 0`. Unlike the XSAVE instruction, `RFBM[2]` does not determine whether XSAVEC saves MXCSR and `MXCSR_MASK`.

## 13.11 OPERATION OF XSAVES

The operation of XSAVES is similar to that of XSAVEC. The main differences are (1) XSAVES can be executed only if `CPL = 0`; (2) XSAVES can operate on the state components whose bits are set in `XCR0 | IA32_XSS` and can thus operate on supervisor state components; and (3) XSAVES uses the modified optimization (see Section 13.6). See Section 13.2 for details of how to determine whether XSAVES is supported.

The XSAVES instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair `EDX:EAX` is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**. `EDX:EAX & (XCR0 | IA32_XSS)` (the logical AND the instruction mask with the logical OR of `XCR0` and `IA32_XSS`) is the **requested-feature bitmap (RFBM)** of the state components to be saved.

The following conditions cause execution of the XSAVES instruction to generate a fault:

- If the XSAVE feature set is not enabled (`CR4.OSXSAVE = 0`), an invalid-opcode exception (`#UD`) occurs.
- If `CR0.TS[bit 3]` is 1, a device-not-available exception (`#NM`) occurs.
- If `CPL > 0` or if the address of the XSAVE area is not 64-byte aligned, a general-protection exception (`#GP`) occurs.<sup>1</sup>

If none of these conditions cause a fault, execution of XSAVES writes the `XSTATE_BV` field of the XSAVE header (see Section 13.4.2), setting `XSTATE_BV[ $i$ ]` ( $0 \leq i \leq 63$ ) as follows:

- If `RFBM[ $i$ ] = 0`, `XSTATE_BV[ $i$ ]` is written as 0.
- If `RFBM[ $i$ ] = 1`, `XSTATE_BV[ $i$ ]` is set to the value of `XINUSE[ $i$ ]` (see below for an exception made for `XSTATE_BV[1]`). Section 13.6 defines `XINUSE` to describe the processor init optimization and specifies the initial configuration of each state component. The nature of that optimization implies the following:
  - If state component  $i$  is in its initial configuration, `XSTATE_BV[ $i$ ]` may be written with either 0 or 1.
  - If state component  $i$  is not in its initial configuration, `XSTATE_BV[ $i$ ]` is written with 1.

`XINUSE[1]` pertains only to the state of the XMM registers and not to MXCSR. However, if `RFBM[1] = 1` and MXCSR does not have the value 1F80H, XSAVES writes `XSTATE_BV[1]` as 1 even if `XINUSE[1] = 0`.

(As explained in Section 13.6, the initial configurations of some state components may depend on whether the processor is in 64-bit mode.)

1. If `CR0.AM = 1`, `CPL = 3`, and `EFLAGS.AC = 1`, an alignment-check exception (`#AC`) may occur instead of `#GP`.

The XSAVES instructions sets bit 63 of the XCOMP\_BV field of the XSAVE header while writing RFBM[62:0] to XCOMP\_BV[62:0]. The XSAVES instruction does not write any part of the XSAVE header other than the XSTATE\_BV and XCOMP\_BV fields.

Execution of XSAVES saves into the XSAVE area those state components corresponding to bits that are set in RFBM (subject to the optimizations described below). State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; the XSAVES instruction always uses the compacted format for the extended region (see Section 13.4.3).

See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes; in particular, see Section 13.5.6 for some special treatment of PT state by XSAVES. See Section 13.13 for details regarding faults caused by memory accesses.

Execution of XSAVES performs the init optimization to reduce the amount of data written to memory. If  $XINUSE[i] = 0$ , state component  $i$  is not saved to the XSAVE area (even if  $RFBM[i] = 1$ ). However, if  $RFBM[1] = 1$  and MXCSR does not have the value 1F80H, XSAVES writes saves all of state component 1 (SSE — including the XMM registers) even if  $XINUSE[1] = 0$ .

Like XSAVEOPT, XSAVES may perform the modified optimization. Each execution of XRSTOR and XRSTORS establishes XRSTOR\_INFO as a 4-tuple  $\langle w, x, y, z \rangle$  (see Section 13.8.3 and Section 13.12). Execution of XSAVES uses the modified optimization only if the following all hold:

- $w = \text{CPL}$ ;
- $x = 1$  if and only if the logical processor is in VMX non-root operation;
- $y$  is the linear address of the XSAVE area being used by XSAVEOPT; and
- $z[63]$  is 1 and  $z[62:0] = \text{RFBM}[62:0]$ . (This last item implies that XSAVES does not use the modified optimization if the last execution of XRSTOR used the standard form and followed the last execution of XRSTORS.)

If XSAVES uses the modified optimization and  $XMODIFIED[i] = 0$  (see Section 13.6), state component  $i$  is not saved to the XSAVE area.

## 13.12 OPERATION OF XRSTORS

The operation of XRSTORS is similar to that of XRSTOR. Three main differences are (1) XRSTORS can be executed only if  $\text{CPL} = 0$ ; (2) XRSTORS can operate on the state components whose bits are set in  $\text{XCR0} \mid \text{IA32\_XSS}$  and can thus operate on supervisor state components; and (3) XRSTORS has only a compacted form (no standard form; see Section 13.8). See Section 13.2 for details of how to determine whether XRSTORS is supported.

The XRSTORS instruction takes a single memory operand, which is an XSAVE area. In addition, the register pair EDX:EAX is an implicit operand used as a state-component bitmap (see Section 13.1) called the **instruction mask**.  $\text{EDX:EAX} \& (\text{XCR0} \mid \text{IA32\_XSS})$  (the logical AND the instruction mask with the logical OR of XCR0 and IA32\_XSS) is the **requested-feature bitmap (RFBM)** of the state components to be restored.

The following conditions cause execution of the XRSTOR instruction to generate a fault:

- If the XSAVE feature set is not enabled ( $\text{CR4.OSXSAVE} = 0$ ), an invalid-opcode exception (#UD) occurs.
- If  $\text{CR0.TS}[\text{bit } 3] = 1$ , a device-not-available exception (#NM) occurs.
- If  $\text{CPL} > 0$  or if the address of the XSAVE area is not 64-byte aligned, a general-protection exception (#GP) occurs.<sup>1</sup>

After checking for these faults, the XRSTORS instruction reads the first 64 bytes of the XSAVE header, including the XSTATE\_BV and XCOMP\_BV fields (see Section 13.4.2). A #GP occurs if any of the following conditions hold for the values read:

- $\text{XCOMP\_BV}[63] = 0$ .
- XCOMP\_BV sets a bit in the range 62:0 that is not set in  $\text{XCR0} \mid \text{IA32\_XSS}$ .
- XSTATE\_BV sets a bit (including bit 63) that is not set in XCOMP\_BV.
- Bytes 63:16 of the XSAVE header are not all 0.

1. If  $\text{CR0.AM} = 1$ ,  $\text{CPL} = 3$ , and  $\text{EFLAGS.AC} = 1$ , an alignment-check exception (#AC) may occur instead of #GP.

If none of these conditions cause a fault, the processor updates each state component  $i$  for which  $RFBM[i] = 1$ . XRSTORS updates state component  $i$  based on the value of bit  $i$  in the XSTATE\_BV field of the XSAVE header:

- If  $XSTATE\_BV[i] = 0$ , the state component is set to its initial configuration. Section 13.6 specifies the initial configuration of each state component. If  $XSTATE\_BV[1] = 0$ , XRSTORS initializes MXCSR to 1F80H.  
State component  $i$  is set to its initial configuration as indicated above if  $RFBM[i] = 1$  and  $XSTATE\_BV[i] = 0$  — even if  $XCOMP\_BV[i] = 0$ . This is true for all values of  $i$ , including 0 (x87 state) and 1 (SSE state).
- If  $XSTATE\_BV[i] = 1$ , the state component is loaded with data from the XSAVE area.<sup>1</sup> See Section 13.5 for specifics for each state component and for details regarding mode-specific operation and operation determined by instruction prefixes; in particular, see Section 13.5.6 for some special treatment of PT state by XRSTORS. See Section 13.13 for details regarding faults caused by memory accesses.

If XRSTORS is restoring a supervisor state component, the instruction causes a general-protection exception (#GP) if it would load any element of that component with an unsupported value (e.g., by setting a reserved bit in an MSR) or if a bit is set in any reserved portion of the state component in the XSAVE area.

State components 0 and 1 are located in the legacy region of the XSAVE area (see Section 13.4.1). Each state component  $i$ ,  $2 \leq i \leq 62$ , is located in the extended region; XRSTORS uses the compacted format for the extended region (see Section 13.4.3).

The MXCSR register is part of SSE state (see Section 13.5.2) and is thus loaded from memory if  $RFBM[1] = XSTATE\_BV[1] = 1$ . XRSTORS causes a general-protection exception (#GP) if it would load MXCSR with an illegal value.

If an execution of XRSTORS causes an exception or a VM exit during or after restoring a supervisor state component, each element of that state component may have the value it held before the XRSTORS execution, the value loaded from the XSAVE area, or the element's initial value (as defined in Section 13.6). See Section 13.5.6 for some special treatment of PT state for the case in which XRSTORS causes an exception or a VM exit.

Like XRSTOR, execution of XRSTORS causes the processor to update its tracking for the init and modified optimizations (see Section 13.6 and Section 13.8.3). The following items provide details:

- The processor updates its tracking for the init optimization as follows:
  - If  $RFBM[i] = 0$ ,  $XINUSE[i]$  is not changed.
  - If  $RFBM[i] = 1$  and  $XSTATE\_BV[i] = 0$ , state component  $i$  may be tracked as init;  $XINUSE[i]$  may be set to 0 or 1.
  - If  $RFBM[i] = 1$  and  $XSTATE\_BV[i] = 1$ , state component  $i$  is tracked as not init;  $XINUSE[i]$  is set to 1.
- The processor updates its tracking for the modified optimization and records information about the XRSTORS execution for future interaction with the XSAVEOPT and XSAVES instructions as follows:
  - If  $RFBM[i] = 0$ , state component  $i$  is tracked as modified;  $XMODIFIED[i]$  is set to 1.
  - If  $RFBM[i] = 1$ , state component  $i$  may be tracked as unmodified;  $XMODIFIED[i]$  may be set to 0 or 1.
  - XRSTOR\_INFO is set to the 4-tuple  $\langle w, x, y, z \rangle$ , where  $w$  is the CPL;  $x$  is 1 if the logical processor is in VMX non-root operation and 0 otherwise;  $y$  is the linear address of the XSAVE area; and  $z$  is  $XCOMP\_BV$  (this implies that  $z[63] = 1$ ).

## 13.13 MEMORY ACCESSSES BY THE XSAVE FEATURE SET

Each instruction in the XSAVE feature set operates on a set of XSAVE-managed state components. The specific set of components on which an instruction operates is determined by the values of XCR0, the IA32\_XSS MSR, EDX:EAX, and (for XRSTOR and XRSTORS) the XSAVE header.

Section 13.4 provides the details necessary to determine the location of each state component for any execution of an instruction in the XSAVE feature set. An execution of an instruction in the XSAVE feature set may access any byte of any state component on which that execution operates.

1. Earlier fault checking ensured that, if the instruction has reached this point in execution and  $XSTATE\_BV[i]$  is 1, then  $XCOMP\_BV[i]$  is also 1.

Section 13.5 provides details of the different XSAVE-managed state components. Some portions of some of these components are accessible only in 64-bit mode. Executions of XRSTOR and XRSTORS outside 64-bit mode will not update those portions; executions of XSAVE, XSAVEC, XSAVEOPT, and XSAVES will not modify the corresponding locations in memory.

Despite this fact, any execution of these instructions outside 64-bit mode may access any byte in any state component on which that execution operates — even those at addresses corresponding to registers that are accessible only in 64-bit mode. As result, such an execution may incur a fault due to an attempt to access such an address.

For example, an execution of XSAVE outside 64-bit mode may incur a page fault if paging does not map as read/write the section of the XSAVE area containing state component 7 (Hi16\_ZMM state) — despite the fact that state component 7 can be accessed only in 64-bit mode.

### 3. Updates to Chapter 1, Volume 2A

Change bars show changes to Chapter 1 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference, A-L*.

-----  
Change to this chapter: Updates to processors covered by manual; added 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series.

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D: Instruction Set Reference* (order numbers 253666, 253667, 326018 and 334569) are part of a set that describes the architecture and programming environment of all Intel 64 and IA-32 architecture processors. Other volumes in this set are:

- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture* (Order Number 253665).
- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D: System Programming Guide* (order numbers 253668, 253669, 326019 and 332831).
- The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers* (order number 335592).

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, describes the basic architecture and programming environment of Intel 64 and IA-32 processors. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D*, describe the instruction set of the processor and the opcode structure. These volumes apply to application programmers and to programmers who write operating systems or executives. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D*, describe the operating-system support environment of Intel 64 and IA-32 processors. These volumes target operating-system and BIOS designers. In addition, the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B*, addresses the programming environment for classes of software that host operating systems. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*, describes the model-specific registers of Intel 64 and IA-32 processors.

## 1.1 INTEL® 64 AND IA-32 PROCESSORS COVERED IN THIS MANUAL

This manual set includes information pertaining primarily to the most recent Intel 64 and IA-32 processors, which include:

- Pentium® processors
- P6 family processors
- Pentium® 4 processors
- Pentium® M processors
- Intel® Xeon® processors
- Pentium® D processors
- Pentium® processor Extreme Editions
- 64-bit Intel® Xeon® processors
- Intel® Core™ Duo processor
- Intel® Core™ Solo processor
- Dual-Core Intel® Xeon® processor LV
- Intel® Core™2 Duo processor
- Intel® Core™2 Quad processor Q6000 series
- Intel® Xeon® processor 3000, 3200 series
- Intel® Xeon® processor 5000 series
- Intel® Xeon® processor 5100, 5300 series
- Intel® Core™2 Extreme processor X7000 and X6800 series
- Intel® Core™2 Extreme processor QX6000 series
- Intel® Xeon® processor 7100 series

## ABOUT THIS MANUAL

- Intel® Pentium® Dual-Core processor
- Intel® Xeon® processor 7200, 7300 series
- Intel® Xeon® processor 5200, 5400, 7400 series
- Intel® Core™2 Extreme processor QX9000 and X9000 series
- Intel® Core™2 Quad processor Q9000 series
- Intel® Core™2 Duo processor E8000, T9000 series
- Intel® Atom™ processor family
- Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are built from 45 nm and 32 nm processes
- Intel® Core™ i7 processor
- Intel® Core™ i5 processor
- Intel® Xeon® processor E7-8800/4800/2800 product families
- Intel® Core™ i7-3930K processor
- 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series
- Intel® Xeon® processor E3-1200 product family
- Intel® Xeon® processor E5-2400/1400 product family
- Intel® Xeon® processor E5-4600/2600/1600 product family
- 3rd generation Intel® Core™ processors
- Intel® Xeon® processor E3-1200 v2 product family
- Intel® Xeon® processor E5-2400/1400 v2 product families
- Intel® Xeon® processor E5-4600/2600/1600 v2 product families
- Intel® Xeon® processor E7-8800/4800/2800 v2 product families
- 4th generation Intel® Core™ processors
- The Intel® Core™ M processor family
- Intel® Core™ i7-59xx Processor Extreme Edition
- Intel® Core™ i7-49xx Processor Extreme Edition
- Intel® Xeon® processor E3-1200 v3 product family
- Intel® Xeon® processor E5-2600/1600 v3 product families
- 5th generation Intel® Core™ processors
- Intel® Xeon® processor D-1500 product family
- Intel® Xeon® processor E5 v4 family
- Intel® Atom™ processor X7-Z8000 and X5-Z8000 series
- Intel® Atom™ processor Z3400 series
- Intel® Atom™ processor Z3500 series
- 6th generation Intel® Core™ processors
- Intel® Xeon® processor E3-1500m v5 product family
- 7th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series
- Intel® Xeon® Processor Scalable Family
- 8th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series

P6 family processors are IA-32 processors based on the P6 family microarchitecture. This includes the Pentium® Pro, Pentium® II, Pentium® III, and Pentium® III Xeon® processors.

The Pentium® 4, Pentium® D, and Pentium® processor Extreme Editions are based on the Intel NetBurst® microarchitecture. Most early Intel® Xeon® processors are based on the Intel NetBurst® microarchitecture. Intel Xeon processor 5000, 7100 series are based on the Intel NetBurst® microarchitecture.

The Intel® Core™ Duo, Intel® Core™ Solo and dual-core Intel® Xeon® processor LV are based on an improved Pentium® M processor microarchitecture.

The Intel® Xeon® processor 3000, 3200, 5100, 5300, 7200, and 7300 series, Intel® Pentium® dual-core, Intel® Core™2 Duo, Intel® Core™2 Quad, and Intel® Core™2 Extreme processors are based on Intel® Core™ microarchitecture.

The Intel® Xeon® processor 5200, 5400, 7400 series, Intel® Core™2 Quad processor Q9000 series, and Intel® Core™2 Extreme processors QX9000, X9000 series, Intel® Core™2 processor E8000 series are based on Enhanced Intel® Core™ microarchitecture.

The Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are based on the Intel® Atom™ microarchitecture and supports Intel 64 architecture.

P6 family, Pentium® M, Intel® Core™ Solo, Intel® Core™ Duo processors, dual-core Intel® Xeon® processor LV, and early generations of Pentium 4 and Intel Xeon processors support IA-32 architecture. The Intel® Atom™ processor Z5xx series support IA-32 architecture.

The Intel® Xeon® processor 3000, 3200, 5000, 5100, 5200, 5300, 5400, 7100, 7200, 7300, 7400 series, Intel® Core™2 Duo, Intel® Core™2 Extreme, Intel® Core™2 Quad processors, Pentium® D processors, Pentium® Dual-Core processor, newer generations of Pentium 4 and Intel Xeon processor family support Intel® 64 architecture.

The Intel® Core™ i7 processor and Intel® Xeon® processor 3400, 5500, 7500 series are based on 45 nm Intel® microarchitecture code name Nehalem. Intel® microarchitecture code name Westmere is a 32 nm version of Intel® microarchitecture code name Nehalem. Intel® Xeon® processor 5600 series, Intel Xeon processor E7 and various Intel Core i7, i5, i3 processors are based on Intel® microarchitecture code name Westmere. These processors support Intel 64 architecture.

The Intel® Xeon® processor E5 family, Intel® Xeon® processor E3-1200 family, Intel® Xeon® processor E7-8800/4800/2800 product families, Intel® Core™ i7-3930K processor, and 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series are based on the Intel® microarchitecture code name Sandy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E7-8800/4800/2800 v2 product families, Intel® Xeon® processor E3-1200 v2 product family and 3rd generation Intel® Core™ processors are based on the Intel® microarchitecture code name Ivy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E5-4600/2600/1600 v2 product families, Intel® Xeon® processor E5-2400/1400 v2 product families and Intel® Core™ i7-49xx Processor Extreme Edition are based on the Intel® microarchitecture code name Ivy Bridge-E and support Intel 64 architecture.

The Intel® Xeon® processor E3-1200 v3 product family and 4th Generation Intel® Core™ processors are based on the Intel® microarchitecture code name Haswell and support Intel 64 architecture.

The Intel® Xeon® processor E5-2600/1600 v3 product families and the Intel® Core™ i7-59xx Processor Extreme Edition are based on the Intel® microarchitecture code name Haswell-E and support Intel 64 architecture.

The Intel® Atom™ processor Z8000 series is based on the Intel microarchitecture code name Airmont.

The Intel® Atom™ processor Z3400 series and the Intel® Atom™ processor Z3500 series are based on the Intel microarchitecture code name Silvermont.

The Intel® Core™ M processor family, 5th generation Intel® Core™ processors, Intel® Xeon® processor D-1500 product family and the Intel® Xeon® processor E5 v4 family are based on the Intel® microarchitecture code name Broadwell and support Intel 64 architecture.

The Intel® Xeon® Processor Scalable Family, Intel® Xeon® processor E3-1500m v5 product family and 6th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Skylake and support Intel 64 architecture.

The 7th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Kaby Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series is based on the Intel® microarchitecture code name Knights Landing and supports Intel 64 architecture.



The 8th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Coffee Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series is based on the Intel® microarchitecture code name Knights Mill and supports Intel 64 architecture.

IA-32 architecture is the instruction set architecture and programming environment for Intel's 32-bit microprocessors. Intel® 64 architecture is the instruction set architecture and programming environment which is the superset of Intel's 32-bit and 64-bit architectures. It is compatible with the IA-32 architecture.

## 1.2 OVERVIEW OF VOLUME 2A, 2B, 2C AND 2D: INSTRUCTION SET REFERENCE

A description of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D* content follows:

**Chapter 1 — About This Manual.** Gives an overview of all seven volumes of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*. It also describes the notational conventions in these manuals and lists related Intel® manuals and documentation of interest to programmers and hardware designers.

**Chapter 2 — Instruction Format.** Describes the machine-level instruction format used for all IA-32 instructions and gives the allowable encodings of prefixes, the operand-identifier byte (ModR/M byte), the addressing-mode specifier byte (SIB byte), and the displacement and immediate bytes.

**Chapter 3 — Instruction Set Reference, A-L.** Describes Intel 64 and IA-32 instructions in detail, including an algorithmic description of operations, the effect on flags, the effect of operand- and address-size attributes, and the exceptions that may be generated. The instructions are arranged in alphabetical order. General-purpose, x87 FPU, Intel MMX™ technology, SSE/SSE2/SSE3/SSSE3/SSE4 extensions, and system instructions are included.

**Chapter 4 — Instruction Set Reference, M-U.** Continues the description of Intel 64 and IA-32 instructions started in Chapter 3. It starts *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B*.

**Chapter 5 — Instruction Set Reference, V-Z.** Continues the description of Intel 64 and IA-32 instructions started in chapters 3 and 4. It provides the balance of the alphabetized list of instructions and starts *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C*.

**Chapter 6— Safer Mode Extensions Reference.** Describes the safer mode extensions (SMX). SMX is intended for a system executive to support launching a measured environment in a platform where the identity of the software controlling the platform hardware can be measured for the purpose of making trust decisions. This chapter starts *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2D*.

**Appendix A — Opcode Map.** Gives an opcode map for the IA-32 instruction set.

**Appendix B — Instruction Formats and Encodings.** Gives the binary encoding of each form of each IA-32 instruction.

**Appendix C — Intel® C/C++ Compiler Intrinsics and Functional Equivalents.** Lists the Intel® C/C++ compiler intrinsics and their assembly code equivalents for each of the IA-32 MMX and SSE/SSE2/SSE3 instructions.

## 1.3 NOTATIONAL CONVENTIONS

This manual uses specific notation for data-structure formats, for symbolic representation of instructions, and for hexadecimal and binary numbers. A review of this notation makes the manual easier to read.

### 1.3.1 Bit and Byte Order

In illustrations of data structures in memory, smaller addresses appear toward the bottom of the figure; addresses increase toward the top. Bit positions are numbered from right to left. The numerical value of a set bit is equal to two raised to the power of the bit position. IA-32 processors are "little endian" machines; this means the bytes of a word are numbered starting from the least significant byte. Figure 1-1 illustrates these conventions.




Figure 1-1. Bit and Byte Order

### 1.3.2 Reserved Bits and Software Compatibility

In many register and memory layout descriptions, certain bits are marked as *reserved*. When bits are marked as reserved, it is essential for compatibility with future processors that software treat these bits as having a future, though unknown, effect. The behavior of reserved bits should be regarded as not only undefined, but unpredictable. Software should follow these guidelines in dealing with reserved bits:

- Do not depend on the states of any reserved bits when testing the values of registers which contain such bits. Mask out the reserved bits before testing.
- Do not depend on the states of any reserved bits when storing to memory or to a register.
- Do not depend on the ability to retain information written into any reserved bits.
- When loading a register, always load the reserved bits with the values indicated in the documentation, if any, or reload them with values previously read from the same register.

#### NOTE

Avoid any software dependence upon the state of reserved bits in IA-32 registers. Depending upon the values of reserved register bits will make software dependent upon the unspecified manner in which the processor handles these bits. Programs that depend upon reserved values risk incompatibility with future processors.

### 1.3.3 Instruction Operands

When instructions are represented symbolically, a subset of the IA-32 assembly language is used. In this subset, an instruction has the following format:

```
label: mnemonic argument1, argument2, argument3
```

where:

- A **label** is an identifier which is followed by a colon.
- A **mnemonic** is a reserved name for a class of instruction opcodes which have the same function.
- The operands *argument1*, *argument2*, and *argument3* are optional. There may be from zero to three operands, depending on the opcode. When present, they take the form of either literals or identifiers for data items. Operand identifiers are either reserved names of registers or are assumed to be assigned to data items declared in another part of the program (which may not be shown in the example).

When two operands are present in an arithmetic or logical instruction, the right operand is the source and the left operand is the destination.

For example:

```
LOADREG: MOV EAX, SUBTOTAL
```

In this example, LOADREG is a label, MOV is the mnemonic identifier of an opcode, EAX is the destination operand, and SUBTOTAL is the source operand. Some assembly languages put the source and destination in reverse order.

### 1.3.4 Hexadecimal and Binary Numbers

Base 16 (hexadecimal) numbers are represented by a string of hexadecimal digits followed by the character H (for example, F82EH). A hexadecimal digit is a character from the following set: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, and F.

Base 2 (binary) numbers are represented by a string of 1s and 0s, sometimes followed by the character B (for example, 1010B). The "B" designation is only used in situations where confusion as to the type of number might arise.

### 1.3.5 Segmented Addressing

The processor uses byte addressing. This means memory is organized and accessed as a sequence of bytes. Whether one or more bytes are being accessed, a byte address is used to locate the byte or bytes in memory. The range of memory that can be addressed is called an address space.

The processor also supports segmented addressing. This is a form of addressing where a program may have many independent address spaces, called segments. For example, a program can keep its code (instructions) and stack in separate segments. Code addresses would always refer to the code space, and stack addresses would always refer to the stack space. The following notation is used to specify a byte address within a segment:

```
Segment-register:Byte-address
```

For example, the following segment address identifies the byte at address FF79H in the segment pointed by the DS register:

```
DS:FF79H
```

The following segment address identifies an instruction address in the code segment. The CS register points to the code segment and the EIP register contains the address of the instruction.

```
CS:EIP
```

### 1.3.6 Exceptions

An exception is an event that typically occurs when an instruction causes an error. For example, an attempt to divide by zero generates an exception. However, some exceptions, such as breakpoints, occur under other conditions. Some types of exceptions may provide error codes. An error code reports additional information about the error. An example of the notation used to show an exception and error code is shown below:

```
#PF(fault code)
```

This example refers to a page-fault exception under conditions where an error code naming a type of fault is reported. Under some conditions, exceptions which produce error codes may not be able to report an accurate code. In this case, the error code is zero, as shown below for a general-protection exception:

```
#GP(0)
```

### 1.3.7 A New Syntax for CPUID, CR, and MSR Values

Obtain feature flags, status, and system information by using the CPUID instruction, by checking control register bits, and by reading model-specific registers. We are moving toward a new syntax to represent this information. See Figure 1-2.

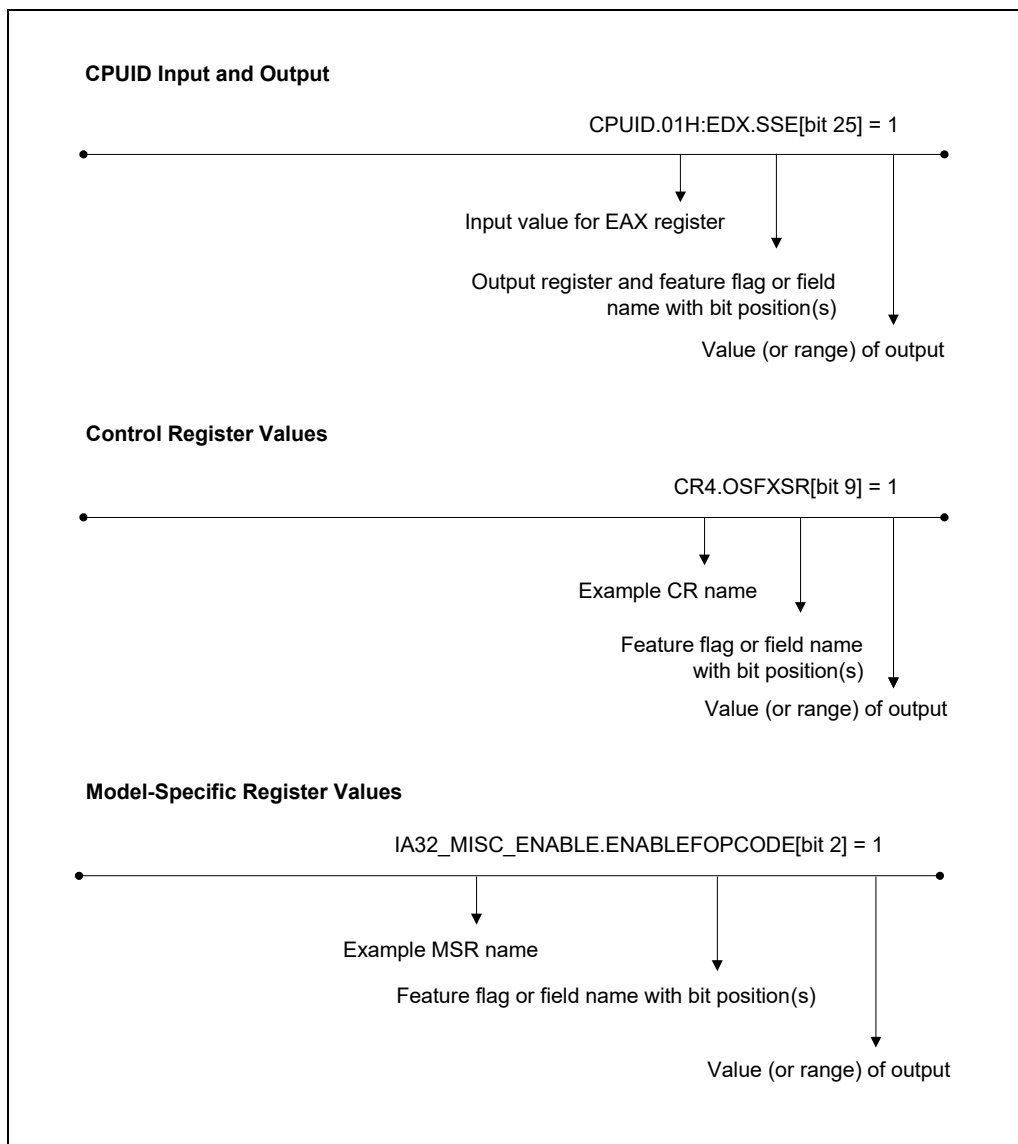


Figure 1-2. Syntax for CPUID, CR, and MSR Data Presentation

## 1.4 RELATED LITERATURE

Literature related to Intel 64 and IA-32 processors is listed and viewable on-line at:

<http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>

See also:

- The data sheet for a particular Intel 64 or IA-32 processor
- The specification update for a particular Intel 64 or IA-32 processor
- Intel® C++ Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Fortran Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Software Development Tools:  
<https://software.intel.com/en-us/intel-sdp-home>

## ABOUT THIS MANUAL

- Intel® 64 and IA-32 Architectures Software Developer's Manual (in one, four or ten volumes):  
<https://software.intel.com/en-us/articles/intel-sdm>
- Intel® 64 and IA-32 Architectures Optimization Reference Manual:  
<https://software.intel.com/en-us/articles/intel-sdm#optimization>
- Intel 64 Architecture x2APIC Specification:  
<http://www.intel.com/content/www/us/en/architecture-and-technology/64-architecture-x2apic-specification.html>
- Intel® Trusted Execution Technology Measured Launched Environment Programming Guide:  
<http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>
- Developing Multi-threaded Applications: A Platform Consistent Approach:  
<https://software.intel.com/sites/default/files/article/147714/51534-developing-multithreaded-applications.pdf>
- Using Spin-Loops on Intel® Pentium® 4 Processor and Intel® Xeon® Processor:  
<https://software.intel.com/sites/default/files/22/30/25602>
- Performance Monitoring Unit Sharing Guide  
<http://software.intel.com/file/30388>

Literature related to selected features in future Intel processors are available at:

- Intel® Architecture Instruction Set Extensions Programming Reference  
<https://software.intel.com/en-us/isa-extensions>
- Intel® Software Guard Extensions (Intel® SGX) Programming Reference  
<https://software.intel.com/en-us/isa-extensions/intel-sgx>

More relevant links are:

- Intel® Developer Zone:  
<https://software.intel.com/en-us>
- Developer centers:  
<http://www.intel.com/content/www/us/en/hardware-developers/developer-centers.html>
- Processor support general link:  
<http://www.intel.com/support/processors/>
- Intel® Hyper-Threading Technology (Intel® HT Technology):  
<http://www.intel.com/technology/platform-technology/hyper-threading/index.htm>

#### 4. Updates to Chapter 3, Volume 2A

Change bars show changes to Chapter 3 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference, A-L*.

-----

Change to this chapter: CPUID instruction updated with new power management details and update to clarify that enumeration of RDPID also Enumerates IA32\_TSC\_AUX.

## CPUID—CPU Identification

| Opcode | Instruction | Op/En | 64-Bit Mode | Compat/Leg Mode | Description   |
|--------|-------------|-------|-------------|-----------------|---|
| 0F A2  | CPUID       | Z0    | Valid       | Valid           | Returns processor identification and feature information to the EAX, EBX, ECX, and EDX registers, as determined by input entered in EAX (in some cases, ECX as well). |

### Instruction Operand Encoding

| Op/En | Operand 1 | Operand 2 | Operand 3 | Operand 4 |
|-------|-----------|-----------|-----------|-----------|
| Z0    | NA        | NA        | NA        | NA        |

### Description

The ID flag (bit 21) in the EFLAGS register indicates support for the CPUID instruction. If a software procedure can set and clear this flag, the processor executing the procedure supports the CPUID instruction. This instruction operates the same in non-64-bit modes and 64-bit mode.

CPUID returns processor identification and feature information in the EAX, EBX, ECX, and EDX registers.<sup>1</sup> The instruction's output is dependent on the contents of the EAX register upon execution (in some cases, ECX as well). For example, the following pseudocode loads EAX with 00H and causes CPUID to return a Maximum Return Value and the Vendor Identification String in the appropriate registers:

```
MOV EAX, 00H
CPUID
```

Table 3-8 shows information returned, depending on the initial value loaded into the EAX register.

Two types of information are returned: basic and extended function information. If a value entered for CPUID.EAX is higher than the maximum input value for basic or extended function for that processor then the data for the highest basic information leaf is returned. For example, using the Intel Core i7 processor, the following is true:

```
CPUID.EAX = 05H (* Returns MONITOR/MWAIT leaf. *)
CPUID.EAX = 0AH (* Returns Architectural Performance Monitoring leaf. *)
CPUID.EAX = 0BH (* Returns Extended Topology Enumeration leaf. *)
CPUID.EAX = 0CH (* INVALID: Returns the same information as CPUID.EAX = 0BH. *)
CPUID.EAX = 80000008H (* Returns linear/physical address size data. *)
CPUID.EAX = 8000000AH (* INVALID: Returns same information as CPUID.EAX = 0BH. *)
```

If a value entered for CPUID.EAX is less than or equal to the maximum input value and the leaf is not supported on that processor then 0 is returned in all the registers.

When CPUID returns the highest basic leaf information as a result of an invalid input EAX value, any dependence on input ECX value in the basic leaf is honored.

CPUID can be executed at any privilege level to serialize instruction execution. Serializing instruction execution guarantees that any modifications to flags, registers, and memory for previous instructions are completed before the next instruction is fetched and executed.

### See also:

“Serializing Instructions” in Chapter 8, “Multiple-Processor Management,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.

“Caching Translation Information” in Chapter 4, “Paging,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.

1. On Intel 64 processors, CPUID clears the high 32 bits of the RAX/RBX/RCX/RDX registers in all modes.

**Table 3-8. Information Returned by CPUID Instruction**

| Initial EAX Value   | Information Provided about the Processor |   |
|---|--|---|
| <i>Basic CPUID Information</i>  |  |   |
| 0H  | EAX                                      | Maximum Input Value for Basic CPUID Information.  |
|   | EBX                                      | "Genu"  |
|   | ECX                                      | "ntel"  |
|   | EDX                                      | "inel"  |
| 01H   | EAX                                      | Version Information: Type, Family, Model, and Stepping ID (see Figure 3-6).   |
|   | EBX                                      | Bits 07 - 00: Brand Index.<br>Bits 15 - 08: CLFLUSH line size (Value * 8 = cache line size in bytes; used also by CLFLUSHOPT).<br>Bits 23 - 16: Maximum number of addressable IDs for logical processors in this physical package*.<br>Bits 31 - 24: Initial APIC ID. |
|   | ECX                                      | Feature Information (see Figure 3-7 and Table 3-10).  |
|   | EDX                                      | Feature Information (see Figure 3-8 and Table 3-11).  |
|   |  | <b>NOTES:</b><br>* The nearest power-of-2 integer that is not smaller than EBX[23:16] is the number of unique initial APIC IDs reserved for addressing different logical processors in a physical package. This field is only valid if CPUID.1.EDX.HTT[bit 28]= 1.    |
| 02H   | EAX                                      | Cache and TLB Information (see Table 3-12).   |
|   | EBX                                      | Cache and TLB Information.  |
|   | ECX                                      | Cache and TLB Information.  |
|   | EDX                                      | Cache and TLB Information.  |
| 03H   | EAX                                      | Reserved.   |
|   | EBX                                      | Reserved.   |
|   | ECX                                      | Bits 00 - 31 of 96 bit processor serial number. (Available in Pentium III processor only; otherwise, the value in this register is reserved.)   |
|   | EDX                                      | Bits 32 - 63 of 96 bit processor serial number. (Available in Pentium III processor only; otherwise, the value in this register is reserved.)   |
|   |  | <b>NOTES:</b><br>Processor serial number (PSN) is not supported in the Pentium 4 processor or later. On all models, use the PSN flag (returned using CPUID) to check for PSN support before accessing the feature.  |
| CPUID leaves above 2 and below 80000000H are visible only when IA32_MISC_ENABLE[bit 22] has its default value of 0. |  |   |
| <i>Deterministic Cache Parameters Leaf</i>  |  |   |
| 04H   |  | <b>NOTES:</b><br>Leaf 04H output depends on the initial value in ECX.*<br>See also: "INPUT EAX = 04H: Returns Deterministic Cache Parameters for Each Level" on page 217.   |
|   | EAX                                      | Bits 04 - 00: Cache Type Field.<br>0 = Null - No more caches.<br>1 = Data Cache.<br>2 = Instruction Cache.<br>3 = Unified Cache.<br>4-31 = Reserved.  |



Table 3-8. Information Returned by CPUID Instruction (Contd.)

| Initial EAX Value | Information Provided about the Processor |   |
|-------------------|--|---|
|                   |  | <p>Bits 07 - 05: Cache Level (starts at 1).<br/>           Bit 08: Self Initializing cache level (does not need SW initialization).<br/>           Bit 09: Fully Associative cache.</p> <p>Bits 13 - 10: Reserved.<br/>           Bits 25 - 14: Maximum number of addressable IDs for logical processors sharing this cache**, ***,<br/>           Bits 31 - 26: Maximum number of addressable IDs for processor cores in the physical package**, ***, ****, *****.</p> <p>EBX Bits 11 - 00: L = System Coherency Line Size**.<br/>           Bits 21 - 12: P = Physical Line partitions**.<br/>           Bits 31 - 22: W = Ways of associativity**.</p> <p>ECX Bits 31-00: S = Number of Sets**.</p> <p>EDX Bit 00: Write-Back Invalidate/Invalidate.<br/>           0 = WBINVD/INVD from threads sharing this cache acts upon lower level caches for threads sharing this cache.<br/>           1 = WBINVD/INVD is not guaranteed to act upon lower level caches of non-originating threads sharing this cache.</p> <p>Bit 01: Cache Inclusiveness.<br/>           0 = Cache is not inclusive of lower cache levels.<br/>           1 = Cache is inclusive of lower cache levels.</p> <p>Bit 02: Complex Cache Indexing.<br/>           0 = Direct mapped cache.<br/>           1 = A complex function is used to index the cache, potentially using all address bits.</p> <p>Bits 31 - 03: Reserved = 0.</p> <p><b>NOTES:</b></p> <p>* If ECX contains an invalid sub leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n+1 is invalid if sub-leaf n returns EAX[4:0] as 0.</p> <p>** Add one to the return value to get the result.</p> <p>***The nearest power-of-2 integer that is not smaller than (1 + EAX[25:14]) is the number of unique initial APIC IDs reserved for addressing different logical processors sharing this cache.</p> <p>**** The nearest power-of-2 integer that is not smaller than (1 + EAX[31:26]) is the number of unique Core_IDs reserved for addressing different processor cores in a physical package. Core ID is a subset of bits of the initial APIC ID.</p> <p>***** The returned value is constant for valid initial values in ECX. Valid ECX values start from 0.</p> |
|                   | <i>MONITOR/MWAIT Leaf</i>                |   |
| 05H               | EAX                                      | Bits 15 - 00: Smallest monitor-line size in bytes (default is processor's monitor granularity).<br>Bits 31 - 16: Reserved = 0.  |
|                   | EBX                                      | Bits 15 - 00: Largest monitor-line size in bytes (default is processor's monitor granularity).<br>Bits 31 - 16: Reserved = 0.   |
|                   | ECX                                      | Bit 00: Enumeration of Monitor-Mwait extensions (beyond EAX and EBX registers) supported.<br>Bit 01: Supports treating interrupts as break-event for MWAIT, even when interrupts disabled.<br>Bits 31 - 02: Reserved.   |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value                        | Information Provided about the Processor |  |
|--|--|--|
|  | EDX                                      | Bits 03 - 00: Number of C0* sub C-states supported using MWAIT.<br>Bits 07 - 04: Number of C1* sub C-states supported using MWAIT.<br>Bits 11 - 08: Number of C2* sub C-states supported using MWAIT.<br>Bits 15 - 12: Number of C3* sub C-states supported using MWAIT.<br>Bits 19 - 16: Number of C4* sub C-states supported using MWAIT.<br>Bits 23 - 20: Number of C5* sub C-states supported using MWAIT.<br>Bits 27 - 24: Number of C6* sub C-states supported using MWAIT.<br>Bits 31 - 28: Number of C7* sub C-states supported using MWAIT.<br><b>NOTE:</b><br>* The definition of C0 through C7 states for MWAIT extension are processor-specific C-states, not ACPI C-states.   |
| <i>Thermal and Power Management Leaf</i> |  |  |
| 06H                                      | EAX                                      | Bit 00: Digital temperature sensor is supported if set.<br>Bit 01: Intel Turbo Boost Technology available (see description of IA32_MISC_ENABLE[38]).<br>Bit 02: ARAT. APIC-Timer-always-running feature is supported if set.<br>Bit 03: Reserved.<br>Bit 04: PLN. Power limit notification controls are supported if set.<br>Bit 05: ECMD. Clock modulation duty cycle extension is supported if set.<br>Bit 06: PTM. Package thermal management is supported if set.<br>Bit 07: HWP. HWP base registers (IA32_PM_ENABLE[bit 0], IA32_HWP_CAPABILITIES, IA32_HWP_REQUEST, IA32_HWP_STATUS) are supported if set.<br>Bit 08: HWP_Notification. IA32_HWP_INTERRUPT MSR is supported if set.<br>Bit 09: HWP_Activity_Window. IA32_HWP_REQUEST[bits 41:32] is supported if set.<br>Bit 10: HWP_Energy_Performance_Preference. IA32_HWP_REQUEST[bits 31:24] is supported if set.<br>Bit 11: HWP_Package_Level_Request. IA32_HWP_REQUEST_PKG MSR is supported if set.<br>Bit 12: Reserved.<br>Bit 13: HDC. HDC base registers IA32_PKG_HDC_CTL, IA32_PM_CTL1, IA32_THREAD_STALL MSRs are supported if set.<br>Bit 14: Intel® Turbo Boost Max Technology 3.0 available.<br>Bit 15: HWP Capabilities. Highest Performance change is supported if set.<br>Bit 16: HWP PECI override is supported if set.<br>Bit 17: Flexible HWP is supported if set.<br>Bit 18: Fast access mode for the IA32_HWP_REQUEST MSR is supported if set.<br>Bit 19: Reserved.<br>Bit 20: Ignoring Idle Logical Processor HWP request is supported if set.<br>Bits 31 - 21: Reserved. |
|  | EBX                                      | Bits 03 - 00: Number of Interrupt Thresholds in Digital Thermal Sensor.<br>Bits 31 - 04: Reserved.   |
|  | ECX                                      | Bit 00: Hardware Coordination Feedback Capability (Presence of IA32_MPERF and IA32_APERF). The capability to provide a measure of delivered processor performance (since last reset of the counters), as a percentage of the expected processor performance when running at the TSC frequency.<br>Bits 02 - 01: Reserved = 0.<br>Bit 03: The processor supports performance-energy bias preference if CPUID.06H:ECX.SETBH[bit 3] is set and it also implies the presence of a new architectural MSR called IA32_ENERGY_PERF_BIAS (1BOH).<br>Bits 31 - 04: Reserved = 0.  |
|  | EDX                                      | Reserved = 0.  |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value   | Information Provided about the Processor  |
|---|---|
| <i>Structured Extended Feature Flags Enumeration Leaf (Output depends on ECX input value)</i> |   |
| 07H   | <p data-bbox="435 338 716 365" style="text-align: center;">Sub-leaf 0 (Input ECX = 0). *</p> <p data-bbox="285 415 1208 443">EAX      Bits 31 - 00: Reports the maximum input value for supported leaf 7 sub-leaves.</p> <p data-bbox="285 457 1442 1415">EBX      Bit 00: FSGSBASE. Supports RDFSBASE/RDGSBASE/WRFSBASE/WRGSBASE if 1.<br/> Bit 01: IA32_TSC_ADJUST MSR is supported if 1.<br/> Bit 02: SGX. Supports Intel® Software Guard Extensions (Intel® SGX Extensions) if 1.<br/> Bit 03: BMI1.<br/> Bit 04: HLE.<br/> Bit 05: AVX2.<br/> Bit 06: FDP_EXCPTN_ONLY. x87 FPU Data Pointer updated only on x87 exceptions if 1.<br/> Bit 07: SMEP. Supports Supervisor-Mode Execution Prevention if 1.<br/> Bit 08: BMI2.<br/> Bit 09: Supports Enhanced REP MOVSB/STOSB if 1.<br/> Bit 10: INVPCID. If 1, supports INVPCID instruction for system software that manages process-context identifiers.<br/> Bit 11: RTM.<br/> Bit 12: RDT-M. Supports Intel® Resource Director Technology (Intel® RDT) Monitoring capability if 1.<br/> Bit 13: Deprecates FPU CS and FPU DS values if 1.<br/> Bit 14: MPX. Supports Intel® Memory Protection Extensions if 1.<br/> Bit 15: RDT-A. Supports Intel® Resource Director Technology (Intel® RDT) Allocation capability if 1.<br/> Bit 16: AVX512F.<br/> Bit 17: AVX512DQ.<br/> Bit 18: RDSEED.<br/> Bit 19: ADX.<br/> Bit 20: SMAP. Supports Supervisor-Mode Access Prevention (and the CLAC/STAC instructions) if 1.<br/> Bit 21: AVX512_IFMA.<br/> Bit 22: Reserved.<br/> Bit 23: CLFLUSHOPT.<br/> Bit 24: CLWB.<br/> Bit 25: Intel Processor Trace.<br/> Bit 26: AVX512PF. (Intel® Xeon Phi™ only.)<br/> Bit 27: AVX512ER. (Intel® Xeon Phi™ only.)<br/> Bit 28: AVX512CD.<br/> Bit 29: SHA. supports Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions) if 1.<br/> Bit 30: AVX512BW.<br/> Bit 31: AVX512VL.</p> <p data-bbox="285 1430 1425 1774">ECX      Bit 00: PREFETCHWT1. (Intel® Xeon Phi™ only.)<br/> Bit 01: AVX512_VBMI.<br/> Bit 02: UMIP. Supports user-mode instruction prevention if 1.<br/> Bit 03: PKU. Supports protection keys for user-mode pages if 1.<br/> Bit 04: OSPKE. If 1, OS has set CR4.PKE to enable protection keys (and the RDPKRU/WRPKRU instructions).<br/> Bits 16 - 5: Reserved.<br/> Bits 21 - 17: The value of MAWAU used by the BNDLDX and BNDSTX instructions in 64-bit mode.<br/> Bit 22: RDPID and IA32_TSC_AUX are available if 1.<br/> Bits 29 - 23: Reserved.<br/> Bit 30: SGX_LC. Supports SGX Launch Configuration if 1.<br/> Bit 31: Reserved.</p> |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value                                | Information Provided about the Processor  |  |
|--|---|--|
|  | EDX   | Reserved.<br><br><b>NOTE:</b><br>* If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n is invalid if n exceeds the value that sub-leaf 0 returns in EAX.   |
| <i>Direct Cache Access Information Leaf</i>      |   |  |
| 09H  | EAX   | Value of bits [31:0] of IA32_PLATFORM_DCA_CAP MSR (address 1F8H).  |
|  | EBX   | Reserved.  |
|  | ECX   | Reserved.  |
|  | EDX   | Reserved.  |
| <i>Architectural Performance Monitoring Leaf</i> |   |  |
| 0AH  | EAX   | Bits 07 - 00: Version ID of architectural performance monitoring.<br>Bits 15 - 08: Number of general-purpose performance monitoring counter per logical processor.<br>Bits 23 - 16: Bit width of general-purpose, performance monitoring counter.<br>Bits 31 - 24: Length of EBX bit vector to enumerate architectural performance monitoring events.  |
|  | EBX   | Bit 00: Core cycle event not available if 1.<br>Bit 01: Instruction retired event not available if 1.<br>Bit 02: Reference cycles event not available if 1.<br>Bit 03: Last-level cache reference event not available if 1.<br>Bit 04: Last-level cache misses event not available if 1.<br>Bit 05: Branch instruction retired event not available if 1.<br>Bit 06: Branch mispredict retired event not available if 1.<br>Bits 31 - 07: Reserved = 0. |
|  | ECX   | Reserved = 0.  |
|  | EDX   | Bits 04 - 00: Number of fixed-function performance counters (if Version ID > 1).<br>Bits 12 - 05: Bit width of fixed-function performance counters (if Version ID > 1).<br>Bits 14 - 13: Reserved = 0.<br>Bit 15: AnyThread deprecation.<br>Bits 31 - 16: Reserved = 0.  |
| <i>Extended Topology Enumeration Leaf</i>        |   |  |
| 0BH  | <b>NOTES:</b><br>Most of Leaf 0BH output depends on the initial value in ECX.<br>The EDX output of leaf 0BH is always valid and does not vary with input value in ECX.<br>Output value in ECX[7:0] always equals input value in ECX[7:0].<br>Sub-leaf index 0 enumerates SMT level. Each subsequent higher sub-leaf index enumerates a higher-level topological entity in hierarchical order.<br>For sub-leaves that return an invalid level-type of 0 in ECX[15:8]; EAX and EBX will return 0.<br>If an input value n in ECX returns the invalid level-type of 0 in ECX[15:8], other input values with ECX > n also return 0 in ECX[15:8]. |  |
|  | EAX   | Bits 04 - 00: Number of bits to shift right on x2APIC ID to get a unique topology ID of the next level type*. All logical processors with the same next level ID share current level.<br>Bits 31 - 05: Reserved.   |
|  | EBX   | Bits 15 - 00: Number of logical processors at this level type. The number reflects configuration as shipped by Intel**.<br>Bits 31 - 16: Reserved.   |

Table 3-8. Information Returned by CPUID Instruction (Contd.)

| Initial EAX Value  | Information Provided about the Processor |  |
|--|--|--|
|  | ECX                                      | Bits 07 - 00: Level number. Same value in ECX input.<br>Bits 15 - 08: Level type***.<br>Bits 31 - 16: Reserved.  |
|  | EDX                                      | Bits 31 - 00: x2APIC ID the current logical processor.   |
|  |  | <p><b>NOTES:</b></p> <p>* Software should use this field (EAX[4:0]) to enumerate processor topology of the system.</p> <p>** Software must not use EBX[15:0] to enumerate processor topology of the system. This value in this field (EBX[15:0]) is only intended for display/diagnostic purposes. The actual number of logical processors available to BIOS/OS/Applications may be different from the value of EBX[15:0], depending on software and platform hardware configurations.</p> <p>*** The value of the "level type" field is not related to level numbers in any way, higher "level type" values do not mean higher levels. Level type field has the following encoding:<br/>0: Invalid.<br/>1: SMT.<br/>2: Core.<br/>3-255: Reserved.</p> |
| <i>Processor Extended State Enumeration Main Leaf (EAX = 0DH, ECX = 0)</i> |  |  |
| 0DH  |  | <p><b>NOTES:</b></p> <p>Leaf 0DH main leaf (ECX = 0).</p>  |
|  | EAX                                      | Bits 31 - 00: Reports the supported bits of the lower 32 bits of XCRO. XCRO[n] can be set to 1 only if EAX[n] is 1.<br>Bit 00: x87 state.<br>Bit 01: SSE state.<br>Bit 02: AVX state.<br>Bits 04 - 03: MPX state.<br>Bits 07 - 05: AVX-512 state.<br>Bit 08: Used for IA32_XSS.<br>Bit 09: PKRU state.<br>Bits 12 - 10: Reserved.<br>Bit 13: Used for IA32_XSS.<br>Bits 31 - 14: Reserved.   |
|  | EBX                                      | Bits 31 - 00: Maximum size (bytes, from the beginning of the XSAVE/XRSTOR save area) required by enabled features in XCRO. May be different than ECX if some features at the end of the XSAVE save area are not enabled.   |
|  | ECX                                      | Bit 31 - 00: Maximum size (bytes, from the beginning of the XSAVE/XRSTOR save area) of the XSAVE/XRSTOR save area required by all supported features in the processor, i.e., all the valid bit fields in XCRO.   |
|  | EDX                                      | Bit 31 - 00: Reports the supported bits of the upper 32 bits of XCRO. XCRO[n+32] can be set to 1 only if EDX[n] is 1.<br>Bits 31 - 00: Reserved.   |
| <i>Processor Extended State Enumeration Sub-leaf (EAX = 0DH, ECX = 1)</i>  |  |  |
| 0DH  | EAX                                      | Bit 00: XSAVEOPT is available.<br>Bit 01: Supports XSAVEC and the compacted form of XRSTOR if set.<br>Bit 02: Supports XGETBV with ECX = 1 if set.<br>Bit 03: Supports XSAVES/XRSTORS and IA32_XSS if set.<br>Bits 31 - 04: Reserved.  |
|  | EBX                                      | Bits 31 - 00: The size in bytes of the XSAVE area containing all states enabled by XCRO   IA32_XSS.  |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value  | Information Provided about the Processor  |
|--|---|
| ECX  | Bits 31 - 00: Reports the supported bits of the lower 32 bits of the IA32_XSS MSR. IA32_XSS[n] can be set to 1 only if ECX[n] is 1.<br>Bits 07 - 00: Used for XCRO.<br>Bit 08: PT state.<br>Bit 09: Used for XCRO.<br>Bits 12 - 10: Reserved.<br>Bit 13: HWP state.<br>Bits 31 - 14: Reserved.  |
| EDX  | Bits 31 - 00: Reports the supported bits of the upper 32 bits of the IA32_XSS MSR. IA32_XSS[n+32] can be set to 1 only if EDX[n] is 1.<br>Bits 31 - 00: Reserved.   |
| <i>Processor Extended State Enumeration Sub-leaves (EAX = 0DH, ECX = n, n &gt; 1)</i>                      |   |
| 0DH  | <p><b>NOTES:</b><br/>                     Leaf 0DH output depends on the initial value in ECX.<br/>                     Each sub-leaf index (starting at position 2) is supported if it corresponds to a supported bit in either the XCRO register or the IA32_XSS MSR.<br/>                     * If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf n (0 ≤ n ≤ 31) is invalid if sub-leaf 0 returns 0 in EAX[n] and sub-leaf 1 returns 0 in ECX[n]. Sub-leaf n (32 ≤ n ≤ 63) is invalid if sub-leaf 0 returns 0 in EDX[n-32] and sub-leaf 1 returns 0 in EDX[n-32].</p> <p>EAX      Bits 31 - 0: The size in bytes (from the offset specified in EBX) of the save area for an extended state feature associated with a valid sub-leaf index, n.</p> <p>EBX      Bits 31 - 0: The offset in bytes of this extended state component's save area from the beginning of the XSAVE/XRSTOR area.<br/>                     This field reports 0 if the sub-leaf index, n, does not map to a valid bit in the XCRO register*.</p> <p>ECX      Bit 00 is set if the bit n (corresponding to the sub-leaf index) is supported in the IA32_XSS MSR; it is clear if bit n is instead supported in XCRO.<br/>                     Bit 01 is set if, when the compacted format of an XSAVE area is used, this extended state component located on the next 64-byte boundary following the preceding state component (otherwise, it is located immediately following the preceding state component).<br/>                     Bits 31 - 02 are reserved.<br/>                     This field reports 0 if the sub-leaf index, n, is invalid*.</p> <p>EDX      This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved.</p> |
| <i>Intel Resource Director Technology (Intel RDT) Monitoring Enumeration Sub-leaf (EAX = 0FH, ECX = 0)</i> |   |
| 0FH  | <p><b>NOTES:</b><br/>                     Leaf 0FH output depends on the initial value in ECX.<br/>                     Sub-leaf index 0 reports valid resource type starting at bit position 1 of EDX.</p> <p>EAX      Reserved.</p> <p>EBX      Bits 31 - 00: Maximum range (zero-based) of RMID within this physical processor of all types.</p> <p>ECX      Reserved.</p> <p>EDX      Bit 00: Reserved.<br/>                     Bit 01: Supports L3 Cache Intel RDT Monitoring if 1.<br/>                     Bits 31 - 02: Reserved.</p>  |
| <i>L3 Cache Intel RDT Monitoring Capability Enumeration Sub-leaf (EAX = 0FH, ECX = 1)</i>                  |   |
| 0FH  | <p><b>NOTES:</b><br/>                     Leaf 0FH output depends on the initial value in ECX.</p> <p>EAX      Reserved.</p>  |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value  | Information Provided about the Processor   |
|--|--|
|  | EBX Bits 31 - 00: Conversion factor from reported IA32_QM_CTR value to occupancy metric (bytes).<br>ECX Maximum range (zero-based) of RMID of this resource type.<br>EDX Bit 00: Supports L3 occupancy monitoring if 1.<br>Bit 01: Supports L3 Total Bandwidth monitoring if 1.<br>Bit 02: Supports L3 Local Bandwidth monitoring if 1.<br>Bits 31 - 03: Reserved.   |
| <i>Intel Resource Director Technology (Intel RDT) Allocation Enumeration Sub-leaf (EAX = 10H, ECX = 0)</i> |  |
| 10H  | <b>NOTES:</b><br>Leaf 10H output depends on the initial value in ECX.<br>Sub-leaf index 0 reports valid resource identification (ResID) starting at bit position 1 of EBX.<br><br>EAX Reserved.<br>EBX Bit 00: Reserved.<br>Bit 01: Supports L3 Cache Allocation Technology if 1.<br>Bit 02: Supports L2 Cache Allocation Technology if 1.<br>Bit 03: Supports Memory Bandwidth Allocation if 1.<br>Bits 31 - 04: Reserved.<br><br>ECX Reserved.<br>EDX Reserved.  |
| <i>L3 Cache Allocation Technology Enumeration Sub-leaf (EAX = 10H, ECX = ResID = 1)</i>                    |  |
| 10H  | <b>NOTES:</b><br>Leaf 10H output depends on the initial value in ECX.<br><br>EAX Bits 04 - 00: Length of the capacity bit mask for the corresponding ResID using minus-one notation.<br>Bits 31 - 05: Reserved.<br><br>EBX Bits 31 - 00: Bit-granular map of isolation/contention of allocation units.<br><br>ECX Bits 01 - 00: Reserved.<br>Bit 02: Code and Data Prioritization Technology supported if 1.<br>Bits 31 - 03: Reserved.<br><br>EDX Bits 15 - 00: Highest COS number supported for this ResID.<br>Bits 31 - 16: Reserved. |
| <i>L2 Cache Allocation Technology Enumeration Sub-leaf (EAX = 10H, ECX = ResID = 2)</i>                    |  |
| 10H  | <b>NOTES:</b><br>Leaf 10H output depends on the initial value in ECX.<br><br>EAX Bits 04 - 00: Length of the capacity bit mask for the corresponding ResID using minus-one notation.<br>Bits 31 - 05: Reserved.<br><br>EBX Bits 31 - 00: Bit-granular map of isolation/contention of allocation units.<br><br>ECX Bits 31 - 00: Reserved.<br><br>EDX Bits 15 - 00: Highest COS number supported for this ResID.<br>Bits 31 - 16: Reserved.   |
| <i>Memory Bandwidth Allocation Enumeration Sub-leaf (EAX = 10H, ECX = ResID = 3)</i>                       |  |
| 10H  | <b>NOTES:</b><br>Leaf 10H output depends on the initial value in ECX.  |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value  | Information Provided about the Processor   |
|--|--|
|  | <p>EAX Bits 11 - 00: Reports the maximum MBA throttling value supported for the corresponding ResID using minus-one notation.<br/>Bits 31 - 12: Reserved.</p> <p>EBX Bits 31 - 00: Reserved.</p> <p>ECX Bits 01 - 00: Reserved.<br/>Bit 02: Reports whether the response of the delay values is linear.<br/>Bits 31 - 03: Reserved.</p> <p>EDX Bits 15 - 00: Highest COS number supported for this ResID.<br/>Bits 31 - 16: Reserved.</p>  |
| <i>Intel SGX Capability Enumeration Leaf, sub-leaf 0 (EAX = 12H, ECX = 0)</i>    |  |
| 12H  | <p><b>NOTES:</b><br/>Leaf 12H sub-leaf 0 (ECX = 0) is supported if CPUID.(EAX=07H, ECX=0H):EBX[SGX] = 1.</p> <p>EAX Bit 00: SGX1. If 1, Indicates Intel SGX supports the collection of SGX1 leaf functions.<br/>Bit 01: SGX2. If 1, Indicates Intel SGX supports the collection of SGX2 leaf functions.<br/>Bits 04 - 02: Reserved.<br/>Bit 05: If 1, indicates Intel SGX supports ENCLV instruction leaves EINCVIRTUAL, EDECVIRTUAL, and ESETCONTEXT.<br/>Bit 06: If 1, indicates Intel SGX supports ENCLS instruction leaves ETRACKC, ERDINFO, ELDBC, and ELDUC.<br/>Bits 31 - 02: Reserved.</p> <p>EBX Bits 31 - 00: MISCSELECT. Bit vector of supported extended SGX features.</p> <p>ECX Bits 31 - 00: Reserved.</p> <p>EDX Bits 07 - 00: MaxEnclaveSize_Not64. The maximum supported enclave size in non-64-bit mode is 2^(EDX[7:0]).<br/>Bits 15 - 08: MaxEnclaveSize_64. The maximum supported enclave size in 64-bit mode is 2^(EDX[15:8]).<br/>Bits 31 - 16: Reserved.</p> |
| <i>Intel SGX Attributes Enumeration Leaf, sub-leaf 1 (EAX = 12H, ECX = 1)</i>    |  |
| 12H  | <p><b>NOTES:</b><br/>Leaf 12H sub-leaf 1 (ECX = 1) is supported if CPUID.(EAX=07H, ECX=0H):EBX[SGX] = 1.</p> <p>EAX Bit 31 - 00: Reports the valid bits of SECS.ATTRIBUTES[31:0] that software can set with ECREATE.</p> <p>EBX Bit 31 - 00: Reports the valid bits of SECS.ATTRIBUTES[63:32] that software can set with ECREATE.</p> <p>ECX Bit 31 - 00: Reports the valid bits of SECS.ATTRIBUTES[95:64] that software can set with ECREATE.</p> <p>EDX Bit 31 - 00: Reports the valid bits of SECS.ATTRIBUTES[127:96] that software can set with ECREATE.</p>   |
| <i>Intel SGX EPC Enumeration Leaf, sub-leaves (EAX = 12H, ECX = 2 or higher)</i> |  |
| 12H  | <p><b>NOTES:</b><br/>Leaf 12H sub-leaf 2 or higher (ECX &gt;= 2) is supported if CPUID.(EAX=07H, ECX=0H):EBX[SGX] = 1.<br/>For sub-leaves (ECX = 2 or higher), definition of EDX,ECX,EBX,EAX[31:4] depends on the sub-leaf type listed below.</p> <p>EAX Bit 03 - 00: Sub-leaf Type<br/>0000b: Indicates this sub-leaf is invalid.<br/>0001b: This sub-leaf enumerates an EPC section. EBX:EAX and EDX:ECX provide information on the Enclave Page Cache (EPC) section.<br/>All other type encodings are reserved.</p> <p>Type 0000b. This sub-leaf is invalid.<br/>EDX:ECX:EBX:EAX return 0.</p>  |



Table 3-8. Information Returned by CPUID Instruction (Contd.)

| Initial EAX Value   | Information Provided about the Processor |   |
|---|--|---|
|   | Type                                     | <p>0001b. This sub-leaf enumerates an EPC sections with EDX:ECX, EBX:EAX defined as follows.</p> <p>EAX[11:04]: Reserved (enumerate 0).<br/>EAX[31:12]: Bits 31:12 of the physical address of the base of the EPC section.</p> <p>EBX[19:00]: Bits 51:32 of the physical address of the base of the EPC section.<br/>EBX[31:20]: Reserved.</p> <p>ECX[03:00]: EPC section property encoding defined as follows:<br/>If EAX[3:0] 0000b, then all bits of the EDX:ECX pair are enumerated as 0.<br/>If EAX[3:0] 0001b, then this section has confidentiality and integrity protection.<br/>All other encodings are reserved.</p> <p>ECX[11:04]: Reserved (enumerate 0).<br/>ECX[31:12]: Bits 31:12 of the size of the corresponding EPC section within the Processor Reserved Memory.</p> <p>EDX[19:00]: Bits 51:32 of the size of the corresponding EPC section within the Processor Reserved Memory.<br/>EDX[31:20]: Reserved.</p>  |
| <i>Intel Processor Trace Enumeration Main Leaf (EAX = 14H, ECX = 0)</i> |  |   |
| 14H   |  | <p><b>NOTES:</b><br/>Leaf 14H main leaf (ECX = 0).</p> <p>EAX Bits 31 - 00: Reports the maximum sub-leaf supported in leaf 14H.</p> <p>EBX Bit 00: If 1, indicates that IA32_RTIT_CTL.CR3Filter can be set to 1, and that IA32_RTIT_CR3_MATCH MSR can be accessed.<br/>Bit 01: If 1, indicates support of Configurable PSB and Cycle-Accurate Mode.<br/>Bit 02: If 1, indicates support of IP Filtering, TraceStop filtering, and preservation of Intel PT MSRs across warm reset.<br/>Bit 03: If 1, indicates support of MTC timing packet and suppression of COFI-based packets.<br/>Bit 04: If 1, indicates support of PTWRITE. Writes can set IA32_RTIT_CTL[12] (PTWEn) and IA32_RTIT_CTL[5] (FUPonPTw), and PTWRITE can generate packets.<br/>Bit 05: If 1, indicates support of Power Event Trace. Writes can set IA32_RTIT_CTL[4] (PwrEvtEn), enabling Power Event Trace packet generation.<br/>Bit 31 - 06: Reserved.</p> <p>ECX Bit 00: If 1, Tracing can be enabled with IA32_RTIT_CTL.ToPA = 1, hence utilizing the ToPA output scheme; IA32_RTIT_OUTPUT_BASE and IA32_RTIT_OUTPUT_MASK_PTRS MSRs can be accessed.<br/>Bit 01: If 1, ToPA tables can hold any number of output entries, up to the maximum allowed by the MaskOffsetTableOffset field of IA32_RTIT_OUTPUT_MASK_PTRS.<br/>Bit 02: If 1, indicates support of Single-Range Output scheme.<br/>Bit 03: If 1, indicates support of output to Trace Transport subsystem.<br/>Bit 30 - 04: Reserved.<br/>Bit 31: If 1, generated packets which contain IP payloads have LIP values, which include the CS base component.</p> <p>EDX Bits 31 - 00: Reserved.</p> |
| <i>Intel Processor Trace Enumeration Sub-leaf (EAX = 14H, ECX = 1)</i>  |  |   |
| 14H   |  | <p>EAX Bits 02 - 00: Number of configurable Address Ranges for filtering.<br/>Bits 15 - 03: Reserved.<br/>Bits 31 - 16: Bitmap of supported MTC period encodings.</p> <p>EBX Bits 15 - 00: Bitmap of supported Cycle Threshold value encodings.<br/>Bit 31 - 16: Bitmap of supported Configurable PSB frequency encodings.</p> <p>ECX Bits 31 - 00: Reserved.</p>   |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value   | Information Provided about the Processor |  |
|---|--|--|
|   | EDX                                      | Bits 31 - 00: Reserved.  |
| <i>Time Stamp Counter and Nominal Core Crystal Clock Information Leaf</i>         |  |  |
| 15H   |  | <p><b>NOTES:</b></p> <p>If EBX[31:0] is 0, the TSC/"core crystal clock" ratio is not enumerated.<br/> EBX[31:0]/EAX[31:0] indicates the ratio of the TSC frequency and the core crystal clock frequency.<br/> If ECX is 0, the nominal core crystal clock frequency is not enumerated.<br/> "TSC frequency" = "core crystal clock frequency" * EBX/EAX.<br/> The core crystal clock may differ from the reference clock, bus clock, or core clock frequencies.</p> <p>EAX     Bits 31 - 00: An unsigned integer which is the denominator of the TSC/"core crystal clock" ratio.<br/> EBX     Bits 31 - 00: An unsigned integer which is the numerator of the TSC/"core crystal clock" ratio.<br/> ECX     Bits 31 - 00: An unsigned integer which is the nominal frequency of the core crystal clock in Hz.<br/> EDX     Bits 31 - 00: Reserved = 0.</p>                           |
| <i>Processor Frequency Information Leaf</i>                                       |  |  |
| 16H   | EAX                                      | Bits 15 - 00: Processor Base Frequency (in MHz).<br>Bits 31 - 16: Reserved = 0.  |
|   | EBX                                      | Bits 15 - 00: Maximum Frequency (in MHz).<br>Bits 31 - 16: Reserved = 0.   |
|   | ECX                                      | Bits 15 - 00: Bus (Reference) Frequency (in MHz).<br>Bits 31 - 16: Reserved = 0.   |
|   | EDX                                      | Reserved.  |
|   |  | <p><b>NOTES:</b></p> <p>* Data is returned from this interface in accordance with the processor's specification and does not reflect actual values. Suitable use of this data includes the display of processor information in like manner to the processor brand string and for determining the appropriate range to use when displaying processor information e.g. frequency history graphs. The returned information should not be used for any other purpose as the returned information does not accurately correlate to information / counters returned by other processor interfaces.</p> <p>While a processor may support the Processor Frequency Information leaf, fields that return a value of zero are not supported.</p>  |
| <i>System-On-Chip Vendor Attribute Enumeration Main Leaf (EAX = 17H, ECX = 0)</i> |  |  |
| 17H   |  | <p><b>NOTES:</b></p> <p>Leaf 17H main leaf (ECX = 0).<br/> Leaf 17H output depends on the initial value in ECX.<br/> Leaf 17H sub-leaves 1 through 3 reports SOC Vendor Brand String.<br/> Leaf 17H is valid if MaxSOCID_Index &gt;= 3.<br/> Leaf 17H sub-leaves 4 and above are reserved.</p> <p>EAX     Bits 31 - 00: MaxSOCID_Index. Reports the maximum input value of supported sub-leaf in leaf 17H.<br/> EBX     Bits 15 - 00: SOC Vendor ID.<br/> Bit 16: IsVendorScheme. If 1, the SOC Vendor ID field is assigned via an industry standard enumeration scheme. Otherwise, the SOC Vendor ID field is assigned by Intel.<br/> Bits 31 - 17: Reserved = 0.<br/> ECX     Bits 31 - 00: Project ID. A unique number an SOC vendor assigns to its SOC projects.<br/> EDX     Bits 31 - 00: Stepping ID. A unique number within an SOC project that an SOC vendor assigns.</p> |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value  | Information Provided about the Processor |   |
|--|--|---|
| <i>System-On-Chip Vendor Attribute Enumeration Sub-leaf (EAX = 17H, ECX = 1..3)</i>                |  |   |
| 17H  | EAX<br>EBX<br>ECX<br>EDX                 | <p>Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string.</p> <p>Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string.</p> <p>Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string.</p> <p>Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string.</p> <p><b>NOTES:</b></p> <p>Leaf 17H output depends on the initial value in ECX.</p> <p>SOC Vendor Brand String is a UTF-8 encoded string padded with trailing bytes of 00H.</p> <p>The complete SOC Vendor Brand String is constructed by concatenating in ascending order of EAX:EBX:ECX:EDX and from the sub-leaf 1 fragment towards sub-leaf 3.</p>   |
| <i>System-On-Chip Vendor Attribute Enumeration Sub-leaves (EAX = 17H, ECX &gt; MaxSOCID_Index)</i> |  |   |
| 17H  | EAX<br>EBX<br>ECX<br>EDX                 | <p><b>NOTES:</b></p> <p>Leaf 17H output depends on the initial value in ECX.</p> <p>Bits 31 - 00: Reserved = 0.</p> <p>Bits 31 - 00: Reserved = 0.</p> <p>Bits 31 - 00: Reserved = 0.</p> <p>Bits 31 - 00: Reserved = 0.</p>  |
| <i>Deterministic Address Translation Parameters Main Leaf (EAX = 18H, ECX = 0)</i>                 |  |   |
| 18H  | EAX<br>EBX<br>ECX                        | <p><b>NOTES:</b></p> <p>Each sub-leaf enumerates a different address translation structure.</p> <p>If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n is invalid if n exceeds the value that sub-leaf 0 returns in EAX. A sub-leaf index is also invalid if EDX[4:0] returns 0. Valid sub-leaves do not need to be contiguous or in any particular order. A valid sub-leaf may be in a higher input ECX value than an invalid sub-leaf or than a valid sub-leaf of a higher or lower-level structure.</p> <p>* Some unified TLBs will allow a single TLB entry to satisfy data read/write and instruction fetches. Others will require separate entries (e.g., one loaded on data read/write and another loaded on an instruction fetch). Please see the <i>Intel® 64 and IA-32 Architectures Optimization Reference Manual</i> for details of a particular product.</p> <p>** Add one to the return value to get the result.</p> <p>Bits 31 - 00: Reports the maximum input value of supported sub-leaf in leaf 18H.</p> <p>Bit 00: 4K page size entries supported by this structure.<br/>           Bit 01: 2MB page size entries supported by this structure.<br/>           Bit 02: 4MB page size entries supported by this structure.<br/>           Bit 03: 1 GB page size entries supported by this structure.<br/>           Bits 07 - 04: Reserved.<br/>           Bits 10 - 08: Partitioning (0: Soft partitioning between the logical processors sharing this structure).<br/>           Bits 15 - 11: Reserved.<br/>           Bits 31 - 16: W = Ways of associativity.</p> <p>Bits 31 - 00: S = Number of Sets.</p> |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value   | Information Provided about the Processor  |
|---|---|
|   | <p>EDX      Bits 04 - 00: Translation cache type field.<br/>                   00000b: Null (indicates this sub-leaf is not valid).<br/>                   00001b: Data TLB.<br/>                   00010b: Instruction TLB.<br/>                   00011b: Unified TLB*.<br/>                   All other encodings are reserved.<br/>           Bits 07 - 05: Translation cache level (starts at 1).<br/>           Bit 08: Fully associative structure.<br/>           Bits 13 - 09: Reserved.<br/>           Bits 25- 14: Maximum number of addressable IDs for logical processors sharing this translation cache**<br/>           Bits 31 - 26: Reserved.</p>  |
| <i>Deterministic Address Translation Parameters Sub-leaf (EAX = 18H, ECX ≥ 1)</i> |   |
| 18H   | <p><b>NOTES:</b><br/>           Each sub-leaf enumerates a different address translation structure.<br/>           If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n is invalid if n exceeds the value that sub-leaf 0 returns in EAX. A sub-leaf index is also invalid if EDX[4:0] returns 0. Valid sub-leaves do not need to be contiguous or in any particular order. A valid sub-leaf may be in a higher input ECX value than an invalid sub-leaf or than a valid sub-leaf of a higher or lower-level structure.<br/>           * Some unified TLBs will allow a single TLB entry to satisfy data read/write and instruction fetches. Others will require separate entries (e.g., one loaded on data read/write and another loaded on an instruction fetch) . Please see the <i>Intel® 64 and IA-32 Architectures Optimization Reference Manual</i> for details of a particular product.<br/>           ** Add one to the return value to get the result.</p> <p>EAX      Bits 31 - 00: Reserved.</p> <p>EBX      Bit 00: 4K page size entries supported by this structure.<br/>                     Bit 01: 2MB page size entries supported by this structure.<br/>                     Bit 02: 4MB page size entries supported by this structure.<br/>                     Bit 03: 1 GB page size entries supported by this structure.<br/>                     Bits 07 - 04: Reserved.<br/>                     Bits 10 - 08: Partitioning (0: Soft partitioning between the logical processors sharing this structure).<br/>                     Bits 15 - 11: Reserved.<br/>                     Bits 31 - 16: W = Ways of associativity.</p> <p>ECX      Bits 31 - 00: S = Number of Sets.</p> <p>EDX      Bits 04 - 00: Translation cache type field.<br/>                     0000b: Null (indicates this sub-leaf is not valid).<br/>                     0001b: Data TLB.<br/>                     0010b: Instruction TLB.<br/>                     0011b: Unified TLB*.<br/>                     All other encodings are reserved.<br/>           Bits 07 - 05: Translation cache level (starts at 1).<br/>           Bit 08: Fully associative structure.<br/>           Bits 13 - 09: Reserved.<br/>           Bits 25- 14: Maximum number of addressable IDs for logical processors sharing this translation cache**<br/>           Bits 31 - 26: Reserved.</p> |
| <i>Unimplemented CPUID Leaf Functions</i>   |   |
| 40000000H<br>-<br>4FFFFFFFH   | Invalid. No existing or future CPU will return processor identification or feature information if the initial EAX value is in the range 40000000H to 4FFFFFFFH.   |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value                          | Information Provided about the Processor |   |
|--|--|---|
| <i>Extended Function CPUID Information</i> |  |   |
| 80000000H                                  | EAX                                      | Maximum Input Value for Extended Function CPUID Information.  |
|  | EBX                                      | Reserved.   |
|  | ECX                                      | Reserved.   |
|  | EDX                                      | Reserved.   |
| 80000001H                                  | EAX                                      | Extended Processor Signature and Feature Bits.  |
|  | EBX                                      | Reserved.   |
|  | ECX                                      | Bit 00: LAHF/SAHF available in 64-bit mode.<br>Bits 04 - 01: Reserved.<br>Bit 05: LZCNT.<br>Bits 07 - 06: Reserved.<br>Bit 08: PREFETCHW.<br>Bits 31 - 09: Reserved.  |
|  | EDX                                      | Bits 10 - 00: Reserved.<br>Bit 11: SYSCALL/SYSRET available in 64-bit mode.<br>Bits 19 - 12: Reserved = 0.<br>Bit 20: Execute Disable Bit available.<br>Bits 25 - 21: Reserved = 0.<br>Bit 26: 1-GByte pages are available if 1.<br>Bit 27: RDTSCP and IA32_TSC_AUX are available if 1.<br>Bit 28: Reserved = 0.<br>Bit 29: Intel® 64 Architecture available if 1.<br>Bits 31 - 30: Reserved = 0. |
| 80000002H                                  | EAX                                      | Processor Brand String.   |
|  | EBX                                      | Processor Brand String Continued.   |
|  | ECX                                      | Processor Brand String Continued.   |
|  | EDX                                      | Processor Brand String Continued.   |
| 80000003H                                  | EAX                                      | Processor Brand String Continued.   |
|  | EBX                                      | Processor Brand String Continued.   |
|  | ECX                                      | Processor Brand String Continued.   |
|  | EDX                                      | Processor Brand String Continued.   |
| 80000004H                                  | EAX                                      | Processor Brand String Continued.   |
|  | EBX                                      | Processor Brand String Continued.   |
|  | ECX                                      | Processor Brand String Continued.   |
|  | EDX                                      | Processor Brand String Continued.   |
| 80000005H                                  | EAX                                      | Reserved = 0.   |
|  | EBX                                      | Reserved = 0.   |
|  | ECX                                      | Reserved = 0.   |
|  | EDX                                      | Reserved = 0.   |
| 80000006H                                  | EAX                                      | Reserved = 0.   |
|  | EBX                                      | Reserved = 0.   |
|  | ECX                                      | Bits 07 - 00: Cache Line size in bytes.<br>Bits 11 - 08: Reserved.<br>Bits 15 - 12: L2 Associativity field *.<br>Bits 31 - 16: Cache size in 1K units.  |
|  | EDX                                      | Reserved = 0.   |

**Table 3-8. Information Returned by CPUID Instruction (Contd.)**

| Initial EAX Value | Information Provided about the Processor  |  |
|-------------------|---|--|
|                   | <p><b>NOTES:</b></p> <p>* L2 associativity field encodings:</p> <p>00H - Disabled.<br/> 01H - Direct mapped.<br/> 02H - 2-way.<br/> 04H - 4-way.<br/> 06H - 8-way.<br/> 08H - 16-way.<br/> 0FH - Fully associative.</p> |  |
| 80000007H         | EAX<br>EBX<br>ECX<br>EDX  | Reserved = 0.<br>Reserved = 0.<br>Reserved = 0.<br>Bits 07 - 00: Reserved = 0.<br>Bit 08: Invariant TSC available if 1.<br>Bits 31 - 09: Reserved = 0.   |
| 80000008H         | EAX<br><br>EBX<br>ECX<br>EDX  | Linear/Physical Address size.<br>Bits 07 - 00: #Physical Address Bits*.<br>Bits 15 - 08: #Linear Address Bits.<br>Bits 31 - 16: Reserved = 0.<br><br>Reserved = 0.<br>Reserved = 0.<br>Reserved = 0. |
|                   | <p><b>NOTES:</b></p> <p>* If CPUID.80000008H:EAX[7:0] is supported, the maximum physical address number supported should come from this field.</p>  |  |

### INPUT EAX = 0: Returns CPUID's Highest Value for Basic Processor Information and the Vendor Identification String

When CPUID executes with EAX set to 0, the processor returns the highest value the CPUID recognizes for returning basic processor information. The value is returned in the EAX register and is processor specific.

A vendor identification string is also returned in EBX, EDX, and ECX. For Intel processors, the string is "GenuineIntel" and is expressed:

EBX ← 756e6547h (\* "Genu", with G in the low eight bits of BL \*)

EDX ← 49656e69h (\* "inel", with i in the low eight bits of DL \*)

ECX ← 6c65746eh (\* "ntel", with n in the low eight bits of CL \*)

### INPUT EAX = 80000000H: Returns CPUID's Highest Value for Extended Processor Information

When CPUID executes with EAX set to 80000000H, the processor returns the highest value the processor recognizes for returning extended processor information. The value is returned in the EAX register and is processor specific.

### IA32\_BIOS\_SIGN\_ID Returns Microcode Update Signature


For processors that support the microcode update facility, the IA32\_BIOS\_SIGN\_ID MSR is loaded with the update signature whenever CPUID executes. The signature is returned in the upper DWORD. For details, see Chapter 9 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.

### INPUT EAX = 01H: Returns Model, Family, Stepping Information

When CPUID executes with EAX set to 01H, version information is returned in EAX (see Figure 3-6). For example: model, family, and processor type for the Intel Xeon processor 5100 series is as follows:

- Model — 1111B
- Family — 0101B
- Processor Type — 00B

See Table 3-9 for available processor type values. Stepping IDs are provided as needed.



**Figure 3-6. Version Information Returned by CPUID in EAX**

**Table 3-9. Processor Type Field**

| Type   | Encoding |
|--|----------|
| Original OEM Processor                                 | 00B      |
| Intel OverDrive™ Processor                             | 01B      |
| Dual processor (not applicable to Intel486 processors) | 10B      |
| Intel reserved   | 11B      |

### NOTE

See Chapter 19 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, for information on identifying earlier IA-32 processors.

The Extended Family ID needs to be examined only when the Family ID is 0FH. Integrate the fields into a display using the following rule:

```
IF Family_ID ≠ 0FH
    THEN DisplayFamily = Family_ID;
    ELSE DisplayFamily = Extended_Family_ID + Family_ID;
    (* Right justify and zero-extend 4-bit field. *)
FI;
(* Show DisplayFamily as HEX field. *)
```

The Extended Model ID needs to be examined only when the Family ID is 06H or 0FH. Integrate the field into a display using the following rule:

```
IF (Family_ID = 06H or Family_ID = 0FH)
    THEN DisplayModel = (Extended_Model_ID << 4) + Model_ID;
    (* Right justify and zero-extend 4-bit field; display Model_ID as HEX field.*)
    ELSE DisplayModel = Model_ID;
FI;
(* Show DisplayModel as HEX field. *)
```

### INPUT EAX = 01H: Returns Additional Information in EBX

When CPUID executes with EAX set to 01H, additional information is returned to the EBX register:

- Brand index (low byte of EBX) — this number provides an entry into a brand string table that contains brand strings for IA-32 processors. More information about this field is provided later in this section.
- CLFLUSH instruction cache line size (second byte of EBX) — this number indicates the size of the cache line flushed by the CLFLUSH and CLFLUSHOPT instructions in 8-byte increments. This field was introduced in the Pentium 4 processor.
- Local APIC ID (high byte of EBX) — this number is the 8-bit ID that is assigned to the local APIC on the processor during power up. This field was introduced in the Pentium 4 processor.

### INPUT EAX = 01H: Returns Feature Information in ECX and EDX

When CPUID executes with EAX set to 01H, feature information is returned in ECX and EDX.

- Figure 3-7 and Table 3-10 show encodings for ECX.
- Figure 3-8 and Table 3-11 show encodings for EDX.

For all feature flags, a 1 indicates that the feature is supported. Use Intel to properly interpret feature flags.

### NOTE

Software must confirm that a processor feature is present using feature flags returned by CPUID prior to using the feature. Software should not depend on future offerings retaining all features.



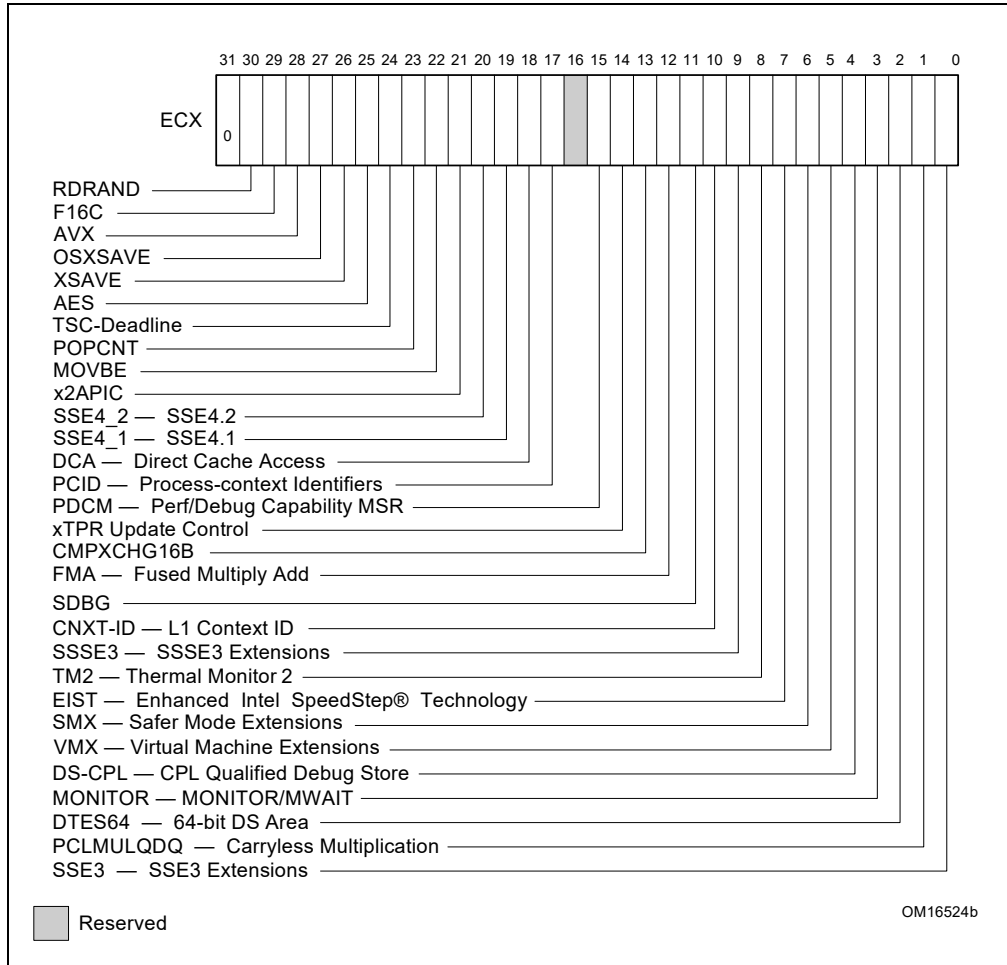



Figure 3-7. Feature Information Returned in the ECX Register

Table 3-10. Feature Information Returned in the ECX Register

| Bit # | Mnemonic  | Description  |
|-------|-----------|--|
| 0     | SSE3      | <b>Streaming SIMD Extensions 3 (SSE3).</b> A value of 1 indicates the processor supports this technology.  |
| 1     | PCLMULQDQ | <b>PCLMULQDQ.</b> A value of 1 indicates the processor supports the PCLMULQDQ instruction.   |
| 2     | DTES64    | <b>64-bit DS Area.</b> A value of 1 indicates the processor supports DS area using 64-bit layout.  |
| 3     | MONITOR   | <b>MONITOR/MWAIT.</b> A value of 1 indicates the processor supports this feature.  |
| 4     | DS-CPL    | <b>CPL Qualified Debug Store.</b> A value of 1 indicates the processor supports the extensions to the Debug Store feature to allow for branch message storage qualified by CPL.  |
| 5     | VMX       | <b>Virtual Machine Extensions.</b> A value of 1 indicates that the processor supports this technology.   |
| 6     | SMX       | <b>Safer Mode Extensions.</b> A value of 1 indicates that the processor supports this technology. See Chapter 6, “Safer Mode Extensions Reference”.                              |
| 7     | EIST      | <b>Enhanced Intel SpeedStep® technology.</b> A value of 1 indicates that the processor supports this technology.   |
| 8     | TM2       | <b>Thermal Monitor 2.</b> A value of 1 indicates whether the processor supports this technology.   |
| 9     | SSSE3     | A value of 1 indicates the presence of the Supplemental Streaming SIMD Extensions 3 (SSSE3). A value of 0 indicates the instruction extensions are not present in the processor. |

Table 3-10. Feature Information Returned in the ECX Register (Contd.)

| Bit # | Mnemonic            | Description  |
|-------|---------------------|--|
| 10    | CNXT-ID             | <b>L1 Context ID.</b> A value of 1 indicates the L1 data cache mode can be set to either adaptive mode or shared mode. A value of 0 indicates this feature is not supported. See definition of the IA32_MISC_ENABLE MSR Bit 24 (L1 Data Cache Context Mode) for details. |
| 11    | SDBG                | A value of 1 indicates the processor supports IA32_DEBUG_INTERFACE MSR for silicon debug.  |
| 12    | FMA                 | A value of 1 indicates the processor supports FMA extensions using YMM state.  |
| 13    | CMPXCHG16B          | <b>CMPXCHG16B Available.</b> A value of 1 indicates that the feature is available. See the “CMPXCHG8B/CMPXCHG16B—Compare and Exchange Bytes” section in this chapter for a description.  |
| 14    | xTPR Update Control | <b>xTPR Update Control.</b> A value of 1 indicates that the processor supports changing IA32_MISC_ENABLE[bit 23].  |
| 15    | PDCM                | <b>Perfmon and Debug Capability:</b> A value of 1 indicates the processor supports the performance and debug feature indication MSR IA32_PERF_CAPABILITIES.  |
| 16    | Reserved            | Reserved   |
| 17    | PCID                | <b>Process-context identifiers.</b> A value of 1 indicates that the processor supports PCIDs and that software may set CR4.PCIDE to 1.   |
| 18    | DCA                 | A value of 1 indicates the processor supports the ability to prefetch data from a memory mapped device.  |
| 19    | SSE4.1              | A value of 1 indicates that the processor supports SSE4.1.   |
| 20    | SSE4.2              | A value of 1 indicates that the processor supports SSE4.2.   |
| 21    | x2APIC              | A value of 1 indicates that the processor supports x2APIC feature.   |
| 22    | MOVBE               | A value of 1 indicates that the processor supports MOVBE instruction.  |
| 23    | POPCNT              | A value of 1 indicates that the processor supports the POPCNT instruction.   |
| 24    | TSC-Deadline        | A value of 1 indicates that the processor’s local APIC timer supports one-shot operation using a TSC deadline value.   |
| 25    | AESNI               | A value of 1 indicates that the processor supports the AESNI instruction extensions.   |
| 26    | XSAVE               | A value of 1 indicates that the processor supports the XSAVE/XRSTOR processor extended states feature, the XSETBV/XGETBV instructions, and XCRO.   |
| 27    | OSXSAVE             | A value of 1 indicates that the OS has set CR4.OSXSAVE[bit 18] to enable XSETBV/XGETBV instructions to access XCRO and to support processor extended state management using XSAVE/XRSTOR.  |
| 28    | AVX                 | A value of 1 indicates the processor supports the AVX instruction extensions.  |
| 29    | F16C                | A value of 1 indicates that processor supports 16-bit floating-point conversion instructions.  |
| 30    | RDRAND              | A value of 1 indicates that processor supports RDRAND instruction.   |
| 31    | Not Used            | Always returns 0.  |



**Figure 3-8. Feature Information Returned in the EDX Register**

**Table 3-11. More on Feature Information Returned in the EDX Register**

| Bit # | Mnemonic | Description  |
|-------|----------|--|
| 0     | FPU      | <b>Floating Point Unit On-Chip.</b> The processor contains an x87 FPU.   |
| 1     | VME      | <b>Virtual 8086 Mode Enhancements.</b> Virtual 8086 mode enhancements, including CR4.VME for controlling the feature, CR4.PVI for protected mode virtual interrupts, software interrupt indirection, expansion of the TSS with the software indirection bitmap, and EFLAGS.VIF and EFLAGS.VIP flags.   |
| 2     | DE       | <b>Debugging Extensions.</b> Support for I/O breakpoints, including CR4.DE for controlling the feature, and optional trapping of accesses to DR4 and DR5.  |
| 3     | PSE      | <b>Page Size Extension.</b> Large pages of size 4 MByte are supported, including CR4.PSE for controlling the feature, the defined dirty bit in PDE (Page Directory Entries), optional reserved bit trapping in CR3, PDEs, and PTEs.  |
| 4     | TSC      | <b>Time Stamp Counter.</b> The RDTSC instruction is supported, including CR4.TSD for controlling privilege.  |
| 5     | MSR      | <b>Model Specific Registers RDMSR and WRMSR Instructions.</b> The RDMSR and WRMSR instructions are supported. Some of the MSRs are implementation dependent.   |
| 6     | PAE      | <b>Physical Address Extension.</b> Physical addresses greater than 32 bits are supported: extended page table entry formats, an extra level in the page translation tables is defined, 2-MByte pages are supported instead of 4 Mbyte pages if PAE bit is 1.   |
| 7     | MCE      | <b>Machine Check Exception.</b> Exception 18 is defined for Machine Checks, including CR4.MCE for controlling the feature. This feature does not define the model-specific implementations of machine-check error logging, reporting, and processor shutdowns. Machine Check exception handlers may have to depend on processor version to do model specific processing of the exception, or test for the presence of the Machine Check feature. |
| 8     | CX8      | <b>CMPXCHG8B Instruction.</b> The compare-and-exchange 8 bytes (64 bits) instruction is supported (implicitly locked and atomic).  |
| 9     | APIC     | <b>APIC On-Chip.</b> The processor contains an Advanced Programmable Interrupt Controller (APIC), responding to memory mapped commands in the physical address range FFFE0000H to FFFE0FFFH (by default - some processors permit the APIC to be relocated).  |
| 10    | Reserved | Reserved   |
| 11    | SEP      | <b>SYSENTER and SYSEXIT Instructions.</b> The SYSENTER and SYSEXIT and associated MSRs are supported.  |
| 12    | MTRR     | <b>Memory Type Range Registers.</b> MTRRs are supported. The MTRRcap MSR contains feature bits that describe what memory types are supported, how many variable MTRRs are supported, and whether fixed MTRRs are supported.  |
| 13    | PGE      | <b>Page Global Bit.</b> The global bit is supported in paging-structure entries that map a page, indicating TLB entries that are common to different processes and need not be flushed. The CR4.PGE bit controls this feature.   |
| 14    | MCA      | <b>Machine Check Architecture.</b> A value of 1 indicates the Machine Check Architecture of reporting machine errors is supported. The MCG_CAP MSR contains feature bits describing how many banks of error reporting MSRs are supported.  |
| 15    | CMOV     | <b>Conditional Move Instructions.</b> The conditional move instruction CMOV is supported. In addition, if x87 FPU is present as indicated by the CPUID.FPU feature bit, then the FCOMI and FCMOV instructions are supported  |
| 16    | PAT      | <b>Page Attribute Table.</b> Page Attribute Table is supported. This feature augments the Memory Type Range Registers (MTRRs), allowing an operating system to specify attributes of memory accessed through a linear address on a 4KB granularity.  |
| 17    | PSE-36   | <b>36-Bit Page Size Extension.</b> 4-MByte pages addressing physical memory beyond 4 GBytes are supported with 32-bit paging. This feature indicates that upper bits of the physical address of a 4-MByte page are encoded in bits 20:13 of the page-directory entry. Such physical addresses are limited by MAXPHYADDR and may be up to 40 bits in size.  |
| 18    | PSN      | <b>Processor Serial Number.</b> The processor supports the 96-bit processor identification number feature and the feature is enabled.  |
| 19    | CLFSH    | <b>CLFLUSH Instruction.</b> CLFLUSH Instruction is supported.  |
| 20    | Reserved | Reserved   |

**Table 3-11. More on Feature Information Returned in the EDX Register (Contd.)**

| Bit # | Mnemonic | Description   |
|-------|----------|---|
| 21    | DS       | <b>Debug Store.</b> The processor supports the ability to write debug information into a memory resident buffer. This feature is used by the branch trace store (BTS) and processor event-based sampling (PEBS) facilities (see Chapter 23, "Introduction to Virtual-Machine Extensions," in the <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C</i> ). |
| 22    | ACPI     | <b>Thermal Monitor and Software Controlled Clock Facilities.</b> The processor implements internal MSRs that allow processor temperature to be monitored and processor performance to be modulated in predefined duty cycles under software control.  |
| 23    | MMX      | <b>Intel MMX Technology.</b> The processor supports the Intel MMX technology.   |
| 24    | FXSR     | <b>FXSAVE and FXRSTOR Instructions.</b> The FXSAVE and FXRSTOR instructions are supported for fast save and restore of the floating point context. Presence of this bit also indicates that CR4.OSFXSR is available for an operating system to indicate that it supports the FXSAVE and FXRSTOR instructions.   |
| 25    | SSE      | <b>SSE.</b> The processor supports the SSE extensions.  |
| 26    | SSE2     | <b>SSE2.</b> The processor supports the SSE2 extensions.  |
| 27    | SS       | <b>Self Snoop.</b> The processor supports the management of conflicting memory types by performing a snoop of its own cache structure for transactions issued to the bus.   |
| 28    | HTT      | <b>Max APIC IDs reserved field is Valid.</b> A value of 0 for HTT indicates there is only a single logical processor in the package and software should assume only a single APIC ID is reserved. A value of 1 for HTT indicates the value in CPUID.1.EBX[23:16] (the Maximum number of addressable IDs for logical processors in this package) is valid for the package.           |
| 29    | TM       | <b>Thermal Monitor.</b> The processor implements the thermal monitor automatic thermal control circuitry (TCC).   |
| 30    | Reserved | Reserved  |
| 31    | PBE      | <b>Pending Break Enable.</b> The processor supports the use of the FERR#/PBE# pin when the processor is in the stop-clock state (STPCLK# is asserted) to signal the processor that an interrupt is pending and that the processor should return to normal operation to handle the interrupt. Bit 10 (PBE enable) in the IA32_MISC_ENABLE MSR enables this capability.               |

**INPUT EAX = 02H: TLB/Cache/Prefetch Information Returned in EAX, EBX, ECX, EDX**

When CPUID executes with EAX set to 02H, the processor returns information about the processor's internal TLBs, cache and prefetch hardware in the EAX, EBX, ECX, and EDX registers. The information is reported in encoded form and fall into the following categories:

- The least-significant byte in register EAX (register AL) will always return 01H. Software should ignore this value and not interpret it as an informational descriptor.
- The most significant bit (bit 31) of each register indicates whether the register contains valid information (set to 0) or is reserved (set to 1).
- If a register contains valid information, the information is contained in 1 byte descriptors. There are four types of encoding values for the byte descriptor, the encoding type is noted in the second column of Table 3-12. Table 3-12 lists the encoding of these descriptors. Note that the order of descriptors in the EAX, EBX, ECX, and EDX registers is not defined; that is, specific bytes are not designated to contain descriptors for specific cache, prefetch, or TLB types. The descriptors may appear in any order. Note also a processor may report a general descriptor type (FFH) and not report any byte descriptor of "cache type" via CPUID leaf 2.

Table 3-12. Encoding of CPUID Leaf 2 Descriptors

| Value | Type    | Description  |
|-------|---------|--|
| 00H   | General | Null descriptor, this byte contains no information   |
| 01H   | TLB     | Instruction TLB: 4 KByte pages, 4-way set associative, 32 entries  |
| 02H   | TLB     | Instruction TLB: 4 MByte pages, fully associative, 2 entries   |
| 03H   | TLB     | Data TLB: 4 KByte pages, 4-way set associative, 64 entries   |
| 04H   | TLB     | Data TLB: 4 MByte pages, 4-way set associative, 8 entries  |
| 05H   | TLB     | Data TLB1: 4 MByte pages, 4-way set associative, 32 entries  |
| 06H   | Cache   | 1st-level instruction cache: 8 KBytes, 4-way set associative, 32 byte line size  |
| 08H   | Cache   | 1st-level instruction cache: 16 KBytes, 4-way set associative, 32 byte line size   |
| 09H   | Cache   | 1st-level instruction cache: 32KBytes, 4-way set associative, 64 byte line size  |
| 0AH   | Cache   | 1st-level data cache: 8 KBytes, 2-way set associative, 32 byte line size   |
| 0BH   | TLB     | Instruction TLB: 4 MByte pages, 4-way set associative, 4 entries   |
| 0CH   | Cache   | 1st-level data cache: 16 KBytes, 4-way set associative, 32 byte line size  |
| 0DH   | Cache   | 1st-level data cache: 16 KBytes, 4-way set associative, 64 byte line size  |
| 0EH   | Cache   | 1st-level data cache: 24 KBytes, 6-way set associative, 64 byte line size  |
| 1DH   | Cache   | 2nd-level cache: 128 KBytes, 2-way set associative, 64 byte line size  |
| 21H   | Cache   | 2nd-level cache: 256 KBytes, 8-way set associative, 64 byte line size  |
| 22H   | Cache   | 3rd-level cache: 512 KBytes, 4-way set associative, 64 byte line size, 2 lines per sector  |
| 23H   | Cache   | 3rd-level cache: 1 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 24H   | Cache   | 2nd-level cache: 1 MBytes, 16-way set associative, 64 byte line size   |
| 25H   | Cache   | 3rd-level cache: 2 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 29H   | Cache   | 3rd-level cache: 4 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 2CH   | Cache   | 1st-level data cache: 32 KBytes, 8-way set associative, 64 byte line size  |
| 30H   | Cache   | 1st-level instruction cache: 32 KBytes, 8-way set associative, 64 byte line size   |
| 40H   | Cache   | No 2nd-level cache or, if processor contains a valid 2nd-level cache, no 3rd-level cache   |
| 41H   | Cache   | 2nd-level cache: 128 KBytes, 4-way set associative, 32 byte line size  |
| 42H   | Cache   | 2nd-level cache: 256 KBytes, 4-way set associative, 32 byte line size  |
| 43H   | Cache   | 2nd-level cache: 512 KBytes, 4-way set associative, 32 byte line size  |
| 44H   | Cache   | 2nd-level cache: 1 MByte, 4-way set associative, 32 byte line size   |
| 45H   | Cache   | 2nd-level cache: 2 MByte, 4-way set associative, 32 byte line size   |
| 46H   | Cache   | 3rd-level cache: 4 MByte, 4-way set associative, 64 byte line size   |
| 47H   | Cache   | 3rd-level cache: 8 MByte, 8-way set associative, 64 byte line size   |
| 48H   | Cache   | 2nd-level cache: 3MByte, 12-way set associative, 64 byte line size   |
| 49H   | Cache   | 3rd-level cache: 4MB, 16-way set associative, 64-byte line size (Intel Xeon processor MP, Family 0FH, Model 06H);<br>2nd-level cache: 4 MByte, 16-way set associative, 64 byte line size |
| 4AH   | Cache   | 3rd-level cache: 6MByte, 12-way set associative, 64 byte line size   |
| 4BH   | Cache   | 3rd-level cache: 8MByte, 16-way set associative, 64 byte line size   |
| 4CH   | Cache   | 3rd-level cache: 12MByte, 12-way set associative, 64 byte line size  |
| 4DH   | Cache   | 3rd-level cache: 16MByte, 16-way set associative, 64 byte line size  |
| 4EH   | Cache   | 2nd-level cache: 6MByte, 24-way set associative, 64 byte line size   |
| 4FH   | TLB     | Instruction TLB: 4 KByte pages, 32 entries   |

Table 3-12. Encoding of CPUID Leaf 2 Descriptors (Contd.)

| Value | Type  | Description   |
|-------|-------|---|
| 50H   | TLB   | Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 64 entries   |
| 51H   | TLB   | Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 128 entries  |
| 52H   | TLB   | Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 256 entries  |
| 55H   | TLB   | Instruction TLB: 2-MByte or 4-MByte pages, fully associative, 7 entries   |
| 56H   | TLB   | Data TLB0: 4 MByte pages, 4-way set associative, 16 entries   |
| 57H   | TLB   | Data TLB0: 4 KByte pages, 4-way associative, 16 entries   |
| 59H   | TLB   | Data TLB0: 4 KByte pages, fully associative, 16 entries   |
| 5AH   | TLB   | Data TLB0: 2 MByte or 4 MByte pages, 4-way set associative, 32 entries  |
| 5BH   | TLB   | Data TLB: 4 KByte and 4 MByte pages, 64 entries   |
| 5CH   | TLB   | Data TLB: 4 KByte and 4 MByte pages, 128 entries  |
| 5DH   | TLB   | Data TLB: 4 KByte and 4 MByte pages, 256 entries  |
| 60H   | Cache | 1st-level data cache: 16 KByte, 8-way set associative, 64 byte line size  |
| 61H   | TLB   | Instruction TLB: 4 KByte pages, fully associative, 48 entries   |
| 63H   | TLB   | Data TLB: 2 MByte or 4 MByte pages, 4-way set associative, 32 entries and a separate array with 1 GByte pages, 4-way set associative, 4 entries |
| 64H   | TLB   | Data TLB: 4 KByte pages, 4-way set associative, 512 entries   |
| 66H   | Cache | 1st-level data cache: 8 KByte, 4-way set associative, 64 byte line size   |
| 67H   | Cache | 1st-level data cache: 16 KByte, 4-way set associative, 64 byte line size  |
| 68H   | Cache | 1st-level data cache: 32 KByte, 4-way set associative, 64 byte line size  |
| 6AH   | Cache | uTLB: 4 KByte pages, 8-way set associative, 64 entries  |
| 6BH   | Cache | DTLB: 4 KByte pages, 8-way set associative, 256 entries   |
| 6CH   | Cache | DTLB: 2M/4M pages, 8-way set associative, 128 entries   |
| 6DH   | Cache | DTLB: 1 GByte pages, fully associative, 16 entries  |
| 70H   | Cache | Trace cache: 12 K- $\mu$ op, 8-way set associative  |
| 71H   | Cache | Trace cache: 16 K- $\mu$ op, 8-way set associative  |
| 72H   | Cache | Trace cache: 32 K- $\mu$ op, 8-way set associative  |
| 76H   | TLB   | Instruction TLB: 2M/4M pages, fully associative, 8 entries  |
| 78H   | Cache | 2nd-level cache: 1 MByte, 4-way set associative, 64byte line size   |
| 79H   | Cache | 2nd-level cache: 128 KByte, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 7AH   | Cache | 2nd-level cache: 256 KByte, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 7BH   | Cache | 2nd-level cache: 512 KByte, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 7CH   | Cache | 2nd-level cache: 1 MByte, 8-way set associative, 64 byte line size, 2 lines per sector  |
| 7DH   | Cache | 2nd-level cache: 2 MByte, 8-way set associative, 64byte line size   |
| 7FH   | Cache | 2nd-level cache: 512 KByte, 2-way set associative, 64-byte line size  |
| 80H   | Cache | 2nd-level cache: 512 KByte, 8-way set associative, 64-byte line size  |
| 82H   | Cache | 2nd-level cache: 256 KByte, 8-way set associative, 32 byte line size  |
| 83H   | Cache | 2nd-level cache: 512 KByte, 8-way set associative, 32 byte line size  |
| 84H   | Cache | 2nd-level cache: 1 MByte, 8-way set associative, 32 byte line size  |
| 85H   | Cache | 2nd-level cache: 2 MByte, 8-way set associative, 32 byte line size  |
| 86H   | Cache | 2nd-level cache: 512 KByte, 4-way set associative, 64 byte line size  |
| 87H   | Cache | 2nd-level cache: 1 MByte, 8-way set associative, 64 byte line size  |

Table 3-12. Encoding of CPUID Leaf 2 Descriptors (Contd.)

| Value | Type     | Description  |
|-------|----------|--|
| A0H   | DTLB     | DTLB: 4k pages, fully associative, 32 entries  |
| B0H   | TLB      | Instruction TLB: 4 KByte pages, 4-way set associative, 128 entries   |
| B1H   | TLB      | Instruction TLB: 2M pages, 4-way, 8 entries or 4M pages, 4-way, 4 entries  |
| B2H   | TLB      | Instruction TLB: 4KByte pages, 4-way set associative, 64 entries   |
| B3H   | TLB      | Data TLB: 4 KByte pages, 4-way set associative, 128 entries  |
| B4H   | TLB      | Data TLB1: 4 KByte pages, 4-way associative, 256 entries   |
| B5H   | TLB      | Instruction TLB: 4KByte pages, 8-way set associative, 64 entries   |
| B6H   | TLB      | Instruction TLB: 4KByte pages, 8-way set associative, 128 entries  |
| BAH   | TLB      | Data TLB1: 4 KByte pages, 4-way associative, 64 entries  |
| C0H   | TLB      | Data TLB: 4 KByte and 4 MByte pages, 4-way associative, 8 entries  |
| C1H   | STLB     | Shared 2nd-Level TLB: 4 KByte/2MByte pages, 8-way associative, 1024 entries  |
| C2H   | DTLB     | DTLB: 4 KByte/2 MByte pages, 4-way associative, 16 entries   |
| C3H   | STLB     | Shared 2nd-Level TLB: 4 KByte /2 MByte pages, 6-way associative, 1536 entries. Also 1GByte pages, 4-way, 16 entries.               |
| C4H   | DTLB     | DTLB: 2M/4M Byte pages, 4-way associative, 32 entries  |
| CAH   | STLB     | Shared 2nd-Level TLB: 4 KByte pages, 4-way associative, 512 entries  |
| D0H   | Cache    | 3rd-level cache: 512 KByte, 4-way set associative, 64 byte line size   |
| D1H   | Cache    | 3rd-level cache: 1 MByte, 4-way set associative, 64 byte line size   |
| D2H   | Cache    | 3rd-level cache: 2 MByte, 4-way set associative, 64 byte line size   |
| D6H   | Cache    | 3rd-level cache: 1 MByte, 8-way set associative, 64 byte line size   |
| D7H   | Cache    | 3rd-level cache: 2 MByte, 8-way set associative, 64 byte line size   |
| D8H   | Cache    | 3rd-level cache: 4 MByte, 8-way set associative, 64 byte line size   |
| DCH   | Cache    | 3rd-level cache: 1.5 MByte, 12-way set associative, 64 byte line size  |
| DDH   | Cache    | 3rd-level cache: 3 MByte, 12-way set associative, 64 byte line size  |
| DEH   | Cache    | 3rd-level cache: 6 MByte, 12-way set associative, 64 byte line size  |
| E2H   | Cache    | 3rd-level cache: 2 MByte, 16-way set associative, 64 byte line size  |
| E3H   | Cache    | 3rd-level cache: 4 MByte, 16-way set associative, 64 byte line size  |
| E4H   | Cache    | 3rd-level cache: 8 MByte, 16-way set associative, 64 byte line size  |
| EAH   | Cache    | 3rd-level cache: 12MByte, 24-way set associative, 64 byte line size  |
| EBH   | Cache    | 3rd-level cache: 18MByte, 24-way set associative, 64 byte line size  |
| ECH   | Cache    | 3rd-level cache: 24MByte, 24-way set associative, 64 byte line size  |
| FOH   | Prefetch | 64-Byte prefetching  |
| F1H   | Prefetch | 128-Byte prefetching   |
| FEH   | General  | CPUID leaf 2 does not report TLB descriptor information; use CPUID leaf 18H to query TLB and other address translation parameters. |
| FFH   | General  | CPUID leaf 2 does not report cache descriptor information, use CPUID leaf 4 to query cache parameters                              |



### Example 3-1. Example of Cache and TLB Interpretation

The first member of the family of Pentium 4 processors returns the following information about caches and TLBs when the CPUID executes with an input value of 2:

```
EAX    66 5B 50 01H
EBX    0H
ECX    0H
EDX    00 7A 70 00H
```

Which means:

- The least-significant byte (byte 0) of register EAX is set to 01H. This value should be ignored.
- The most-significant bit of all four registers (EAX, EBX, ECX, and EDX) is set to 0, indicating that each register contains valid 1-byte descriptors.
- Bytes 1, 2, and 3 of register EAX indicate that the processor has:
  - 50H - a 64-entry instruction TLB, for mapping 4-KByte and 2-MByte or 4-MByte pages.
  - 5BH - a 64-entry data TLB, for mapping 4-KByte and 4-MByte pages.
  - 66H - an 8-KByte 1st level data cache, 4-way set associative, with a 64-Byte cache line size.
- The descriptors in registers EBX and ECX are valid, but contain NULL descriptors.
- Bytes 0, 1, 2, and 3 of register EDX indicate that the processor has:
  - 00H - NULL descriptor.
  - 70H - Trace cache: 12 K- $\mu$ op, 8-way set associative.
  - 7AH - a 256-KByte 2nd level cache, 8-way set associative, with a sectored, 64-byte cache line size.
  - 00H - NULL descriptor.

### INPUT EAX = 04H: Returns Deterministic Cache Parameters for Each Level

When CPUID executes with EAX set to 04H and ECX contains an index value, the processor returns encoded data that describe a set of deterministic cache parameters (for the cache level associated with the input in ECX). Valid index values start from 0.

Software can enumerate the deterministic cache parameters for each level of the cache hierarchy starting with an index value of 0, until the parameters report the value associated with the cache type field is 0. The architecturally defined fields reported by deterministic cache parameters are documented in Table 3-8.

This Cache Size in Bytes

$$= (\text{Ways} + 1) * (\text{Partitions} + 1) * (\text{Line\_Size} + 1) * (\text{Sets} + 1)$$

$$= (\text{EBX}[31:22] + 1) * (\text{EBX}[21:12] + 1) * (\text{EBX}[11:0] + 1) * (\text{ECX} + 1)$$

The CPUID leaf 04H also reports data that can be used to derive the topology of processor cores in a physical package. This information is constant for all valid index values. Software can query the raw data reported by executing CPUID with EAX=04H and ECX=0 and use it as part of the topology enumeration algorithm described in Chapter 8, “Multiple-Processor Management,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.

### INPUT EAX = 05H: Returns MONITOR and MWAIT Features

When CPUID executes with EAX set to 05H, the processor returns information about features available to MONITOR/MWAIT instructions. The MONITOR instruction is used for address-range monitoring in conjunction with MWAIT instruction. The MWAIT instruction optionally provides additional extensions for advanced power management. See Table 3-8.

### INPUT EAX = 06H: Returns Thermal and Power Management Features

When CPUID executes with EAX set to 06H, the processor returns information about thermal and power management features. See Table 3-8.

**INPUT EAX = 07H: Returns Structured Extended Feature Enumeration Information**

When CPUID executes with EAX set to 07H and ECX = 0, the processor returns information about the maximum input value for sub-leaves that contain extended feature flags. See Table 3-8.

When CPUID executes with EAX set to 07H and the input value of ECX is invalid (see leaf 07H entry in Table 3-8), the processor returns 0 in EAX/EBX/ECX/EDX. In subleaf 0, EAX returns the maximum input value of the highest leaf 7 sub-leaf, and EBX, ECX & EDX contain information of extended feature flags.

**INPUT EAX = 09H: Returns Direct Cache Access Information**

When CPUID executes with EAX set to 09H, the processor returns information about Direct Cache Access capabilities. See Table 3-8.

**INPUT EAX = 0AH: Returns Architectural Performance Monitoring Features**

When CPUID executes with EAX set to 0AH, the processor returns information about support for architectural performance monitoring capabilities. Architectural performance monitoring is supported if the version ID (see Table 3-8) is greater than Pn 0. See Table 3-8.

For each version of architectural performance monitoring capability, software must enumerate this leaf to discover the programming facilities and the architectural performance events available in the processor. The details are described in Chapter 23, “Introduction to Virtual-Machine Extensions,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3C*.

**INPUT EAX = 0BH: Returns Extended Topology Information**

When CPUID executes with EAX set to 0BH, the processor returns information about extended topology enumeration data. Software must detect the presence of CPUID leaf 0BH by verifying (a) the highest leaf index supported by CPUID is  $\geq 0BH$ , and (b) CPUID.0BH:EBX[15:0] reports a non-zero value. See Table 3-8.

**INPUT EAX = 0DH: Returns Processor Extended States Enumeration Information**

When CPUID executes with EAX set to 0DH and ECX = 0, the processor returns information about the bit-vector representation of all processor state extensions that are supported in the processor and storage size requirements of the XSAVE/XRSTOR area. See Table 3-8.

When CPUID executes with EAX set to 0DH and ECX = n (n > 1, and is a valid sub-leaf index), the processor returns information about the size and offset of each processor extended state save area within the XSAVE/XRSTOR area. See Table 3-8. Software can use the forward-extendable technique depicted below to query the valid sub-leaves and obtain size and offset information for each processor extended state save area:

```
For i = 2 to 62 // sub-leaf 1 is reserved
  IF (CPUID.(EAX=0DH, ECX=0):VECTOR[i] = 1) // VECTOR is the 64-bit value of EDX:EAX
    Execute CPUID.(EAX=0DH, ECX = i) to examine size and offset for sub-leaf i;
  FI;
```

**INPUT EAX = 0FH: Returns Intel Resource Director Technology (Intel RDT) Monitoring Enumeration Information**

When CPUID executes with EAX set to 0FH and ECX = 0, the processor returns information about the bit-vector representation of QoS monitoring resource types that are supported in the processor and maximum range of RMID values the processor can use to monitor of any supported resource types. Each bit, starting from bit 1, corresponds to a specific resource type if the bit is set. The bit position corresponds to the sub-leaf index (or ResID) that software must use to query QoS monitoring capability available for that type. See Table 3-8.

When CPUID executes with EAX set to 0FH and ECX = n (n  $\geq$  1, and is a valid ResID), the processor returns information software can use to program IA32\_PQR\_ASSOC, IA32\_QM\_EVTSEL MSRs before reading QoS data from the IA32\_QM\_CTR MSR.

**INPUT EAX = 10H: Returns Intel Resource Director Technology (Intel RDT) Allocation Enumeration Information**

When CPUID executes with EAX set to 10H and ECX = 0, the processor returns information about the bit-vector representation of QoS Enforcement resource types that are supported in the processor. Each bit, starting from bit

1, corresponds to a specific resource type if the bit is set. The bit position corresponds to the sub-leaf index (or ResID) that software must use to query QoS enforcement capability available for that type. See Table 3-8.

When CPUID executes with EAX set to 10H and ECX = n (n >= 1, and is a valid ResID), the processor returns information about available classes of service and range of QoS mask MSRs that software can use to configure each class of services using capability bit masks in the QoS Mask registers, IA32\_resourceType\_Mask\_n.

#### **INPUT EAX = 12H: Returns Intel SGX Enumeration Information**

When CPUID executes with EAX set to 12H and ECX = 0H, the processor returns information about Intel SGX capabilities. See Table 3-8.

When CPUID executes with EAX set to 12H and ECX = 1H, the processor returns information about Intel SGX attributes. See Table 3-8.

When CPUID executes with EAX set to 12H and ECX = n (n > 1), the processor returns information about Intel SGX Enclave Page Cache. See Table 3-8.

#### **INPUT EAX = 14H: Returns Intel Processor Trace Enumeration Information**

When CPUID executes with EAX set to 14H and ECX = 0H, the processor returns information about Intel Processor Trace extensions. See Table 3-8.

When CPUID executes with EAX set to 14H and ECX = n (n > 0 and less than the number of non-zero bits in CPUID.(EAX=14H, ECX= 0H).EAX), the processor returns information about packet generation in Intel Processor Trace. See Table 3-8.

#### **INPUT EAX = 15H: Returns Time Stamp Counter and Nominal Core Crystal Clock Information**

When CPUID executes with EAX set to 15H and ECX = 0H, the processor returns information about Time Stamp Counter and Core Crystal Clock. See Table 3-8.

#### **INPUT EAX = 16H: Returns Processor Frequency Information**

When CPUID executes with EAX set to 16H, the processor returns information about Processor Frequency Information. See Table 3-8.

#### **INPUT EAX = 17H: Returns System-On-Chip Information**

When CPUID executes with EAX set to 17H, the processor returns information about the System-On-Chip Vendor Attribute Enumeration. See Table 3-8.

#### **INPUT EAX = 18H: Returns Deterministic Address Translation Parameters Information**

When CPUID executes with EAX set to 18H, the processor returns information about the Deterministic Address Translation Parameters. See Table 3-8.

### **METHODS FOR RETURNING BRANDING INFORMATION**

Use the following techniques to access branding information:


1. Processor brand string method.
2. Processor brand index; this method uses a software supplied brand string table.

These two methods are discussed in the following sections. For methods that are available in early processors, see Section: "Identification of Earlier IA-32 Processors" in Chapter 19 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*.

#### **The Processor Brand String Method**

Figure 3-9 describes the algorithm used for detection of the brand string. Processor brand identification software should execute this algorithm on all Intel 64 and IA-32 processors.

This method (introduced with Pentium 4 processors) returns an ASCII brand identification string and the Processor Base frequency of the processor to the EAX, EBX, ECX, and EDX registers.



**Figure 3-9. Determination of Support for the Processor Brand String**

**How Brand Strings Work**

To use the brand string method, execute CPUID with EAX input of 8000002H through 80000004H. For each input value, CPUID returns 16 ASCII characters using EAX, EBX, ECX, and EDX. The returned string will be NULL-terminated.


Table 3-13 shows the brand string that is returned by the first processor in the Pentium 4 processor family.

**Table 3-13. Processor Brand String Returned with Pentium 4 Processor**

| EAX Input Value | Return Values  | ASCII Equivalent                      |
|-----------------|--|---------------------------------------|
| 80000002H       | EAX = 20202020H<br>EBX = 20202020H<br>ECX = 20202020H<br>EDX = 6E492020H | " " "<br>" " "<br>" " "<br>"nl "      |
| 80000003H       | EAX = 286C6574H<br>EBX = 50202952H<br>ECX = 69746E65H<br>EDX = 52286D75H | "(let"<br>"P )R"<br>"itne"<br>"R(mu"  |
| 80000004H       | EAX = 20342029H<br>EBX = 20555043H<br>ECX = 30303531H<br>EDX = 007A484DH | " 4 )"<br>" UPC"<br>"0051"<br>"\0zHM" |

### Extracting the Processor Frequency from Brand Strings

Figure 3-10 provides an algorithm which software can use to extract the Processor Base frequency from the processor brand string.



**Figure 3-10. Algorithm for Extracting Processor Frequency**

## The Processor Brand Index Method

The brand index method (introduced with Pentium® III Xeon® processors) provides an entry point into a brand identification table that is maintained in memory by system software and is accessible from system- and user-level code. In this table, each brand index is associate with an ASCII brand identification string that identifies the official Intel family and model number of a processor.

When CPUID executes with EAX set to 1, the processor returns a brand index to the low byte in EBX. Software can then use this index to locate the brand identification string for the processor in the brand identification table. The first entry (brand index 0) in this table is reserved, allowing for backward compatibility with processors that do not support the brand identification feature. Starting with processor signature family ID = 0FH, model = 03H, brand index method is no longer supported. Use brand string method instead.

Table 3-14 shows brand indices that have identification strings associated with them.

**Table 3-14. Mapping of Brand Indices; and Intel 64 and IA-32 Processor Brand Strings**

| Brand Index | Brand String  |
|-------------|---|
| 00H         | This processor does not support the brand identification feature  |
| 01H         | Intel(R) Celeron(R) processor <sup>1</sup>  |
| 02H         | Intel(R) Pentium(R) III processor <sup>1</sup>  |
| 03H         | Intel(R) Pentium(R) III Xeon(R) processor; If processor signature = 000006B1h, then Intel(R) Celeron(R) processor |
| 04H         | Intel(R) Pentium(R) III processor   |
| 06H         | Mobile Intel(R) Pentium(R) III processor-M  |
| 07H         | Mobile Intel(R) Celeron(R) processor <sup>1</sup>   |
| 08H         | Intel(R) Pentium(R) 4 processor   |
| 09H         | Intel(R) Pentium(R) 4 processor   |
| 0AH         | Intel(R) Celeron(R) processor <sup>1</sup>  |
| 0BH         | Intel(R) Xeon(R) processor; If processor signature = 00000F13h, then Intel(R) Xeon(R) processor MP                |
| 0CH         | Intel(R) Xeon(R) processor MP   |
| 0EH         | Mobile Intel(R) Pentium(R) 4 processor-M; If processor signature = 00000F13h, then Intel(R) Xeon(R) processor     |
| 0FH         | Mobile Intel(R) Celeron(R) processor <sup>1</sup>   |
| 11H         | Mobile Genuine Intel(R) processor   |
| 12H         | Intel(R) Celeron(R) M processor   |
| 13H         | Mobile Intel(R) Celeron(R) processor <sup>1</sup>   |
| 14H         | Intel(R) Celeron(R) processor   |
| 15H         | Mobile Genuine Intel(R) processor   |
| 16H         | Intel(R) Pentium(R) M processor   |
| 17H         | Mobile Intel(R) Celeron(R) processor <sup>1</sup>   |
| 18H - 0FFH  | RESERVED  |

### NOTES:

1. Indicates versions of these processors that were introduced after the Pentium III

## IA-32 Architecture Compatibility

CPUID is not supported in early models of the Intel486 processor or in any IA-32 processor earlier than the Intel486 processor.

## Operation

IA32\_BIOS\_SIGN\_ID MSR ← Update with installed microcode revision number;

CASE (EAX) OF

EAX = 0:

EAX ← Highest basic function input value understood by CPUID;

EBX ← Vendor identification string;

EDX ← Vendor identification string;

ECX ← Vendor identification string;

BREAK;

EAX = 1H:

EAX[3:0] ← Stepping ID;

EAX[7:4] ← Model;

EAX[11:8] ← Family;

EAX[13:12] ← Processor type;

EAX[15:14] ← Reserved;

EAX[19:16] ← Extended Model;

EAX[27:20] ← Extended Family;

EAX[31:28] ← Reserved;

EBX[7:0] ← Brand Index; (\* Reserved if the value is zero. \*)

EBX[15:8] ← CLFLUSH Line Size;

EBX[16:23] ← Reserved; (\* Number of threads enabled = 2 if MT enable fuse set. \*)

EBX[24:31] ← Initial APIC ID;

ECX ← Feature flags; (\* See Figure 3-7. \*)

EDX ← Feature flags; (\* See Figure 3-8. \*)

BREAK;

EAX = 2H:

EAX ← Cache and TLB information;

EBX ← Cache and TLB information;

ECX ← Cache and TLB information;

EDX ← Cache and TLB information;

BREAK;

EAX = 3H:

EAX ← Reserved;

EBX ← Reserved;

ECX ← ProcessorSerialNumber[31:0];

(\* Pentium III processors only, otherwise reserved. \*)

EDX ← ProcessorSerialNumber[63:32];

(\* Pentium III processors only, otherwise reserved. \*)

BREAK

EAX = 4H:

EAX ← Deterministic Cache Parameters Leaf; (\* See Table 3-8. \*)

EBX ← Deterministic Cache Parameters Leaf;

ECX ← Deterministic Cache Parameters Leaf;

EDX ← Deterministic Cache Parameters Leaf;

BREAK;

EAX = 5H:

EAX ← MONITOR/MWAIT Leaf; (\* See Table 3-8. \*)

EBX ← MONITOR/MWAIT Leaf;

ECX ← MONITOR/MWAIT Leaf;

EDX ← MONITOR/MWAIT Leaf;

BREAK;

EAX = 6H:

EAX ← Thermal and Power Management Leaf; (\* See Table 3-8. \*)

EBX ← Thermal and Power Management Leaf;

ECX ← Thermal and Power Management Leaf;

EDX ← Thermal and Power Management Leaf;

BREAK;

EAX = 7H:

EAX ← Structured Extended Feature Flags Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Structured Extended Feature Flags Enumeration Leaf;

ECX ← Structured Extended Feature Flags Enumeration Leaf;

EDX ← Structured Extended Feature Flags Enumeration Leaf;

BREAK;

EAX = 8H:

EAX ← Reserved = 0;

EBX ← Reserved = 0;

ECX ← Reserved = 0;

EDX ← Reserved = 0;

BREAK;

EAX = 9H:

EAX ← Direct Cache Access Information Leaf; (\* See Table 3-8. \*)

EBX ← Direct Cache Access Information Leaf;

ECX ← Direct Cache Access Information Leaf;

EDX ← Direct Cache Access Information Leaf;

BREAK;

EAX = AH:

EAX ← Architectural Performance Monitoring Leaf; (\* See Table 3-8. \*)

EBX ← Architectural Performance Monitoring Leaf;

ECX ← Architectural Performance Monitoring Leaf;

EDX ← Architectural Performance Monitoring Leaf;

BREAK

EAX = BH:

EAX ← Extended Topology Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Extended Topology Enumeration Leaf;

ECX ← Extended Topology Enumeration Leaf;

EDX ← Extended Topology Enumeration Leaf;

BREAK;

EAX = CH:

EAX ← Reserved = 0;

EBX ← Reserved = 0;

ECX ← Reserved = 0;

EDX ← Reserved = 0;

BREAK;

EAX = DH:

EAX ← Processor Extended State Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Processor Extended State Enumeration Leaf;

ECX ← Processor Extended State Enumeration Leaf;

EDX ← Processor Extended State Enumeration Leaf;

BREAK;

EAX = EH:

EAX ← Reserved = 0;

EBX ← Reserved = 0;

ECX ← Reserved = 0;

EDX ← Reserved = 0;

BREAK;



EAX = FH:

EAX ← Intel Resource Director Technology Monitoring Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Intel Resource Director Technology Monitoring Enumeration Leaf;

ECX ← Intel Resource Director Technology Monitoring Enumeration Leaf;

EDX ← Intel Resource Director Technology Monitoring Enumeration Leaf;

BREAK;

EAX = 10H:

EAX ← Intel Resource Director Technology Allocation Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Intel Resource Director Technology Allocation Enumeration Leaf;

ECX ← Intel Resource Director Technology Allocation Enumeration Leaf;

EDX ← Intel Resource Director Technology Allocation Enumeration Leaf;

BREAK;

EAX = 12H:

EAX ← Intel SGX Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Intel SGX Enumeration Leaf;

ECX ← Intel SGX Enumeration Leaf;

EDX ← Intel SGX Enumeration Leaf;

BREAK;

EAX = 14H:

EAX ← Intel Processor Trace Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Intel Processor Trace Enumeration Leaf;

ECX ← Intel Processor Trace Enumeration Leaf;

EDX ← Intel Processor Trace Enumeration Leaf;

BREAK;

EAX = 15H:

EAX ← Time Stamp Counter and Nominal Core Crystal Clock Information Leaf; (\* See Table 3-8. \*)

EBX ← Time Stamp Counter and Nominal Core Crystal Clock Information Leaf;

ECX ← Time Stamp Counter and Nominal Core Crystal Clock Information Leaf;

EDX ← Time Stamp Counter and Nominal Core Crystal Clock Information Leaf;

BREAK;

EAX = 16H:

EAX ← Processor Frequency Information Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Processor Frequency Information Enumeration Leaf;

ECX ← Processor Frequency Information Enumeration Leaf;

EDX ← Processor Frequency Information Enumeration Leaf;

BREAK;

EAX = 17H:

EAX ← System-On-Chip Vendor Attribute Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← System-On-Chip Vendor Attribute Enumeration Leaf;

ECX ← System-On-Chip Vendor Attribute Enumeration Leaf;

EDX ← System-On-Chip Vendor Attribute Enumeration Leaf;

BREAK;

EAX = 18H:

EAX ← Deterministic Address Translation Parameters Enumeration Leaf; (\* See Table 3-8. \*)

EBX ← Deterministic Address Translation Parameters Enumeration Leaf;

ECX ← Deterministic Address Translation Parameters Enumeration Leaf;

EDX ← Deterministic Address Translation Parameters Enumeration Leaf;

BREAK;

EAX = 80000000H:

EAX ← Highest extended function input value understood by CPUID;

EBX ← Reserved;

ECX ← Reserved;

EDX ← Reserved;

BREAK;

EAX = 80000001H:

EAX ← Reserved;  
 EBX ← Reserved;  
 ECX ← Extended Feature Bits (\* See Table 3-8.\*);  
 EDX ← Extended Feature Bits (\* See Table 3-8. \*);

BREAK;

EAX = 80000002H:

EAX ← Processor Brand String;  
 EBX ← Processor Brand String, continued;  
 ECX ← Processor Brand String, continued;  
 EDX ← Processor Brand String, continued;

BREAK;

EAX = 80000003H:

EAX ← Processor Brand String, continued;  
 EBX ← Processor Brand String, continued;  
 ECX ← Processor Brand String, continued;  
 EDX ← Processor Brand String, continued;

BREAK;

EAX = 80000004H:

EAX ← Processor Brand String, continued;  
 EBX ← Processor Brand String, continued;  
 ECX ← Processor Brand String, continued;  
 EDX ← Processor Brand String, continued;

BREAK;

EAX = 80000005H:

EAX ← Reserved = 0;  
 EBX ← Reserved = 0;  
 ECX ← Reserved = 0;  
 EDX ← Reserved = 0;

BREAK;

EAX = 80000006H:

EAX ← Reserved = 0;  
 EBX ← Reserved = 0;  
 ECX ← Cache information;  
 EDX ← Reserved = 0;

BREAK;

EAX = 80000007H:

EAX ← Reserved = 0;  
 EBX ← Reserved = 0;  
 ECX ← Reserved = 0;  
 EDX ← Reserved = Misc Feature Flags;

BREAK;

EAX = 80000008H:

EAX ← Reserved = Physical Address Size Information;  
 EBX ← Reserved = Virtual Address Size Information;  
 ECX ← Reserved = 0;  
 EDX ← Reserved = 0;

BREAK;

EAX >= 40000000H and EAX <= 4FFFFFFFH:

DEFAULT: (\* EAX = Value outside of recognized range for CPUID. \*)

(\* If the highest basic information leaf data depend on ECX input value, ECX is honored.\*)

EAX ← Reserved; (\* Information returned for highest basic information leaf. \*)  
 EBX ← Reserved; (\* Information returned for highest basic information leaf. \*)  
 ECX ← Reserved; (\* Information returned for highest basic information leaf. \*)

EDX ← Reserved; (\* Information returned for highest basic information leaf. \*)  
BREAK;  
ESAC;

### Flags Affected

None.

### Exceptions (All Operating Modes)

#UD                    If the LOCK prefix is used.  
In earlier IA-32 processors that do not support the CPUID instruction, execution of the instruction results in an invalid opcode (#UD) exception being generated.

## 5. Updates to Chapter 4, Volume 2B

Change bars show changes to Chapter 4 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference, M-U*.

-----  
Changes to this chapter: Updates to MOVUPS, PACKUSDW, PTWRITE and SQRTPD instructions.

## MOVUPS—Move Unaligned Packed Single-Precision Floating-Point Values

| Opcode/<br>Instruction                                  | Op / En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|---|---------|------------------------------|--------------------------|---|
| NP OF 10 /r<br>MOVUPS xmm1, xmm2/m128                   | A       | V/V                          | SSE                      | Move unaligned packed single-precision floating-point from xmm2/mem to xmm1.                            |
| NP OF 11 /r<br>MOVUPS xmm2/m128, xmm1                   | B       | V/V                          | SSE                      | Move unaligned packed single-precision floating-point from xmm1 to xmm2/mem.                            |
| VEX.128.OF.WIG 10 /r<br>VMOVUPS xmm1, xmm2/m128         | A       | V/V                          | AVX                      | Move unaligned packed single-precision floating-point from xmm2/mem to xmm1.                            |
| VEX.128.OF.WIG 11 /r<br>VMOVUPS xmm2/m128, xmm1         | B       | V/V                          | AVX                      | Move unaligned packed single-precision floating-point from xmm1 to xmm2/mem.                            |
| VEX.256.OF.WIG 10 /r<br>VMOVUPS ymm1, ymm2/m256         | A       | V/V                          | AVX                      | Move unaligned packed single-precision floating-point from ymm2/mem to ymm1.                            |
| VEX.256.OF.WIG 11 /r<br>VMOVUPS ymm2/m256, ymm1         | B       | V/V                          | AVX                      | Move unaligned packed single-precision floating-point from ymm1 to ymm2/mem.                            |
| EVEX.128.OF.W0 10 /r<br>VMOVUPS xmm1 {k1}{z}, xmm2/m128 | C       | V/V                          | AVX512VL<br>AVX512F      | Move unaligned packed single-precision floating-point values from xmm2/m128 to xmm1 using writemask k1. |
| EVEX.256.OF.W0 10 /r<br>VMOVUPS ymm1 {k1}{z}, ymm2/m256 | C       | V/V                          | AVX512VL<br>AVX512F      | Move unaligned packed single-precision floating-point values from ymm2/m256 to ymm1 using writemask k1. |
| EVEX.512.OF.W0 10 /r<br>VMOVUPS zmm1 {k1}{z}, zmm2/m512 | C       | V/V                          | AVX512F                  | Move unaligned packed single-precision floating-point values from zmm2/m512 to zmm1 using writemask k1. |
| EVEX.128.OF.W0 11 /r<br>VMOVUPS xmm2/m128 {k1}{z}, xmm1 | D       | V/V                          | AVX512VL<br>AVX512F      | Move unaligned packed single-precision floating-point values from xmm1 to xmm2/m128 using writemask k1. |
| EVEX.256.OF.W0 11 /r<br>VMOVUPS ymm2/m256 {k1}{z}, ymm1 | D       | V/V                          | AVX512VL<br>AVX512F      | Move unaligned packed single-precision floating-point values from ymm1 to ymm2/m256 using writemask k1. |
| EVEX.512.OF.W0 11 /r<br>VMOVUPS zmm2/m512 {k1}{z}, zmm1 | D       | V/V                          | AVX512F                  | Move unaligned packed single-precision floating-point values from zmm1 to zmm2/m512 using writemask k1. |

### Instruction Operand Encoding

| Op/En | Tuple Type | Operand 1     | Operand 2     | Operand 3 | Operand 4 |
|-------|------------|---------------|---------------|-----------|-----------|
| A     | NA         | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |
| B     | NA         | ModRM:r/m (w) | ModRM:reg (r) | NA        | NA        |
| C     | Full Mem   | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |
| D     | Full Mem   | ModRM:r/m (w) | ModRM:reg (r) | NA        | NA        |

### Description

Note: VEX.vvvv and EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD.

#### EVEX.512 encoded version:

Moves 512 bits of packed single-precision floating-point values from the source operand (second operand) to the destination operand (first operand). This instruction can be used to load a ZMM register from a 512-bit float32 memory location, to store the contents of a ZMM register into memory. The destination operand is updated according to the writemask.

**VEX.256 and EVEX.256 encoded versions:**

Moves 256 bits of packed single-precision floating-point values from the source operand (second operand) to the destination operand (first operand). This instruction can be used to load a YMM register from a 256-bit memory location, to store the contents of a YMM register into a 256-bit memory location, or to move data between two YMM registers. Bits (MAXVL-1:256) of the destination register are zeroed.

128-bit versions:

Moves 128 bits of packed single-precision floating-point values from the source operand (second operand) to the destination operand (first operand). This instruction can be used to load an XMM register from a 128-bit memory location, to store the contents of an XMM register into a 128-bit memory location, or to move data between two XMM registers.

**128-bit Legacy SSE version:** Bits (MAXVL-1:128) of the corresponding destination register remain unchanged.

When the source or destination operand is a memory operand, the operand may be unaligned without causing a general-protection exception (#GP) to be generated.

**VEX.128 and EVEX.128 encoded versions:** Bits (MAXVL-1:128) of the destination register are zeroed.

**Operation****VMOVUPS (EVEX encoded versions, register-copy form)**

(KL, VL) = (4, 128), (8, 256), (16, 512)

FOR j ← 0 TO KL-1

  i ← j \* 32

  IF k1[j] OR \*no writemask\*

    THEN DEST[i+31:i] ← SRC[i+31:i]

  ELSE

    IF \*merging-masking\* ; merging-masking

      THEN \*DEST[i+31:i] remains unchanged\*

    ELSE DEST[i+31:i] ← 0 ; zeroing-masking

  FI

  FI;

ENDFOR

DEST[MAXVL-1:VL] ← 0

**VMOVUPS (EVEX encoded versions, store-form)**

(KL, VL) = (4, 128), (8, 256), (16, 512)

FOR j ← 0 TO KL-1

  i ← j \* 32

  IF k1[j] OR \*no writemask\*

    THEN DEST[i+31:i] ← SRC[i+31:i]

  ELSE \*DEST[i+31:i] remains unchanged\* ; merging-masking

  FI;

ENDFOR;

**VMOVUPS (EVEX encoded versions, load-form)**

(KL, VL) = (4, 128), (8, 256), (16, 512)

FOR j ← 0 TO KL-1

i ← j \* 32

IF k1[j] OR \*no writemask\*

THEN DEST[i+31:i] ← SRC[i+31:i]

ELSE

IF \*merging-masking\* ; merging-masking

THEN \*DEST[i+31:i] remains unchanged\*

ELSE DEST[i+31:i] ← 0 ; zeroing-masking

FI

FI;

ENDFOR

DEST[MAXVL-1:VL] ← 0

**VMOVUPS (VEX.256 encoded version, load - and register copy)**

DEST[255:0] ← SRC[255:0]

DEST[MAXVL-1:256] ← 0

**VMOVUPS (VEX.256 encoded version, store-form)**

DEST[255:0] ← SRC[255:0]

**VMOVUPS (VEX.128 encoded version)**

DEST[127:0] ← SRC[127:0]

DEST[MAXVL-1:128] ← 0

**MOVUPS (128-bit load- and register-copy- form Legacy SSE version)**

DEST[127:0] ← SRC[127:0]

DEST[MAXVL-1:128] (Unmodified)

**(V)MOVUPS (128-bit store-form version)**

DEST[127:0] ← SRC[127:0]

**Intel C/C++ Compiler Intrinsic Equivalent**

VMOVUPS \_\_m512 \_\_mm512\_loadu\_ps( void \* s);

VMOVUPS \_\_m512 \_\_mm512\_mask\_loadu\_ps(\_\_m512 a, \_\_mmask16 k, void \* s);

VMOVUPS \_\_m512 \_\_mm512\_maskz\_loadu\_ps( \_\_mmask16 k, void \* s);

VMOVUPS void \_\_mm512\_storeu\_ps( void \* d, \_\_m512 a);

VMOVUPS void \_\_mm512\_mask\_storeu\_ps( void \* d, \_\_mmask8 k, \_\_m512 a);

VMOVUPS \_\_m256 \_\_mm256\_mask\_loadu\_ps(\_\_m256 a, \_\_mmask8 k, void \* s);

VMOVUPS \_\_m256 \_\_mm256\_maskz\_loadu\_ps( \_\_mmask8 k, void \* s);

VMOVUPS void \_\_mm256\_mask\_storeu\_ps( void \* d, \_\_mmask8 k, \_\_m256 a);

VMOVUPS \_\_m128 \_\_mm\_mask\_loadu\_ps(\_\_m128 a, \_\_mmask8 k, void \* s);

VMOVUPS \_\_m128 \_\_mm\_maskz\_loadu\_ps( \_\_mmask8 k, void \* s);

VMOVUPS void \_\_mm\_mask\_storeu\_ps( void \* d, \_\_mmask8 k, \_\_m128 a);

MOVUPS \_\_m256 \_\_mm256\_loadu\_ps ( float \* p);

MOVUPS void \_\_mm256\_storeu\_ps( float \*p, \_\_m256 a);

MOVUPS \_\_m128 \_\_mm\_loadu\_ps ( float \* p);

MOVUPS void \_\_mm\_storeu\_ps( float \*p, \_\_m128 a);

**SIMD Floating-Point Exceptions**

None

**Other Exceptions**

Non-EVEX-encoded instruction, see Exceptions Type 4.

Note treatment of #AC varies;

EVEX-encoded instruction, see Exceptions Type E4.nb.

#UD                    If EVEX.vvvv != 1111B or VEX.vvvv != 1111B.



## PACKUSDW—Pack with Unsigned Saturation

| Opcode/<br>Instruction   | Op /<br>En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|--|------------|------------------------------|--------------------------|--|
| 66 0F 38 2B /r<br>PACKUSDW <i>xmm1, xmm2/m128</i>  | A          | V/V                          | SSE4_1                   | Convert 4 packed signed doubleword integers from <i>xmm1</i> and 4 packed signed doubleword integers from <i>xmm2/m128</i> into 8 packed unsigned word integers in <i>xmm1</i> using unsigned saturation.                              |
| VEX.NDS.128.66.0F38 2B /r<br>VPACKUSDW <i>xmm1,xmm2,<br/>xmm3/m128</i>                     | B          | V/V                          | AVX                      | Convert 4 packed signed doubleword integers from <i>xmm2</i> and 4 packed signed doubleword integers from <i>xmm3/m128</i> into 8 packed unsigned word integers in <i>xmm1</i> using unsigned saturation.                              |
| VEX.NDS.256.66.0F38 2B /r<br>VPACKUSDW <i>ymm1, ymm2,<br/>ymm3/m256</i>                    | B          | V/V                          | AVX2                     | Convert 8 packed signed doubleword integers from <i>ymm2</i> and 8 packed signed doubleword integers from <i>ymm3/m256</i> into 16 packed unsigned word integers in <i>ymm1</i> using unsigned saturation.                             |
| EVEX.NDS.128.66.0F38.W0 2B /r<br>VPACKUSDW <i>xmm1{k1}{z},<br/>xmm2, xmm3/m128/m32bcst</i> | C          | V/V                          | AVX512VL<br>AVX512BW     | Convert packed signed doubleword integers from <i>xmm2</i> and packed signed doubleword integers from <i>xmm3/m128/m32bcst</i> into packed unsigned word integers in <i>xmm1</i> using unsigned saturation under writemask <i>k1</i> . |
| EVEX.NDS.256.66.0F38.W0 2B /r<br>VPACKUSDW <i>ymm1{k1}{z},<br/>ymm2, ymm3/m256/m32bcst</i> | C          | V/V                          | AVX512VL<br>AVX512BW     | Convert packed signed doubleword integers from <i>ymm2</i> and packed signed doubleword integers from <i>ymm3/m256/m32bcst</i> into packed unsigned word integers in <i>ymm1</i> using unsigned saturation under writemask <i>k1</i> . |
| EVEX.NDS.512.66.0F38.W0 2B /r<br>VPACKUSDW <i>zmm1{k1}{z},<br/>zmm2, zmm3/m512/m32bcst</i> | C          | V/V                          | AVX512BW                 | Convert packed signed doubleword integers from <i>zmm2</i> and packed signed doubleword integers from <i>zmm3/m512/m32bcst</i> into packed unsigned word integers in <i>zmm1</i> using unsigned saturation under writemask <i>k1</i> . |

### Instruction Operand Encoding

| Op/En | Tuple Type | Operand 1        | Operand 2     | Operand 3     | Operand 4 |
|-------|------------|------------------|---------------|---------------|-----------|
| A     | NA         | ModRM:reg (r, w) | ModRM:r/m (r) | NA            | NA        |
| B     | NA         | ModRM:reg (w)    | VEX.vvvv (r)  | ModRM:r/m (r) | NA        |
| C     | Full       | ModRM:reg (w)    | EVEX.vvvv (r) | ModRM:r/m (r) | NA        |

### Description

Converts packed signed doubleword integers in the first and second source operands into packed unsigned word integers using unsigned saturation to handle overflow conditions. If the signed doubleword value is beyond the range of an unsigned word (that is, greater than FFFFH or less than 0000H), the saturated unsigned word integer value of FFFFH or 0000H, respectively, is stored in the destination.

EVEX encoded versions: The first source operand is a ZMM/YMM/XMM register. The second source operand is a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 32-bit memory location. The destination operand is a ZMM register, updated conditionally under the writemask *k1*.

VEX.256 encoded version: The first source operand is a YMM register. The second source operand is a YMM register or a 256-bit memory location. The destination operand is a YMM register. The upper bits (MAXVL-1:256) of the corresponding ZMM register destination are zeroed.

VEX.128 encoded version: The first source operand is an XMM register. The second source operand is an XMM register or 128-bit memory location. The destination operand is an XMM register. The upper bits (MAXVL-1:128) of the corresponding ZMM register destination are zeroed.

128-bit Legacy SSE version: The first source operand is an XMM register. The second operand can be an XMM register or an 128-bit memory location. The destination is not distinct from the first source XMM register and the upper bits (MAXVL-1:128) of the corresponding destination register destination are unmodified.

**Operation****PACKUSDW (Legacy SSE instruction)**

$TMP[15:0] \leftarrow (DEST[31:0] < 0) ? 0 : DEST[15:0];$   
 $DEST[15:0] \leftarrow (DEST[31:0] > FFFFH) ? FFFFH : TMP[15:0];$   
 $TMP[31:16] \leftarrow (DEST[63:32] < 0) ? 0 : DEST[47:32];$   
 $DEST[31:16] \leftarrow (DEST[63:32] > FFFFH) ? FFFFH : TMP[31:16];$   
 $TMP[47:32] \leftarrow (DEST[95:64] < 0) ? 0 : DEST[79:64];$   
 $DEST[47:32] \leftarrow (DEST[95:64] > FFFFH) ? FFFFH : TMP[47:32];$   
 $TMP[63:48] \leftarrow (DEST[127:96] < 0) ? 0 : DEST[111:96];$   
 $DEST[63:48] \leftarrow (DEST[127:96] > FFFFH) ? FFFFH : TMP[63:48];$   
 $TMP[79:64] \leftarrow (SRC[31:0] < 0) ? 0 : SRC[15:0];$   
 $DEST[79:64] \leftarrow (SRC[31:0] > FFFFH) ? FFFFH : TMP[79:64];$   
 $TMP[95:80] \leftarrow (SRC[63:32] < 0) ? 0 : SRC[47:32];$   
 $DEST[95:80] \leftarrow (SRC[63:32] > FFFFH) ? FFFFH : TMP[95:80];$   
 $TMP[111:96] \leftarrow (SRC[95:64] < 0) ? 0 : SRC[79:64];$   
 $DEST[111:96] \leftarrow (SRC[95:64] > FFFFH) ? FFFFH : TMP[111:96];$   
 $TMP[127:112] \leftarrow (SRC[127:96] < 0) ? 0 : SRC[111:96];$   
 $DEST[127:112] \leftarrow (SRC[127:96] > FFFFH) ? FFFFH : TMP[127:112];$   
 $DEST[MAXVL-1:128]$  (Unmodified)

**PACKUSDW (VEX.128 encoded version)**

$TMP[15:0] \leftarrow (SRC1[31:0] < 0) ? 0 : SRC1[15:0];$   
 $DEST[15:0] \leftarrow (SRC1[31:0] > FFFFH) ? FFFFH : TMP[15:0];$   
 $TMP[31:16] \leftarrow (SRC1[63:32] < 0) ? 0 : SRC1[47:32];$   
 $DEST[31:16] \leftarrow (SRC1[63:32] > FFFFH) ? FFFFH : TMP[31:16];$   
 $TMP[47:32] \leftarrow (SRC1[95:64] < 0) ? 0 : SRC1[79:64];$   
 $DEST[47:32] \leftarrow (SRC1[95:64] > FFFFH) ? FFFFH : TMP[47:32];$   
 $TMP[63:48] \leftarrow (SRC1[127:96] < 0) ? 0 : SRC1[111:96];$   
 $DEST[63:48] \leftarrow (SRC1[127:96] > FFFFH) ? FFFFH : TMP[63:48];$   
 $TMP[79:64] \leftarrow (SRC2[31:0] < 0) ? 0 : SRC2[15:0];$   
 $DEST[79:64] \leftarrow (SRC2[31:0] > FFFFH) ? FFFFH : TMP[79:64];$   
 $TMP[95:80] \leftarrow (SRC2[63:32] < 0) ? 0 : SRC2[47:32];$   
 $DEST[95:80] \leftarrow (SRC2[63:32] > FFFFH) ? FFFFH : TMP[95:80];$   
 $TMP[111:96] \leftarrow (SRC2[95:64] < 0) ? 0 : SRC2[79:64];$   
 $DEST[111:96] \leftarrow (SRC2[95:64] > FFFFH) ? FFFFH : TMP[111:96];$   
 $TMP[127:112] \leftarrow (SRC2[127:96] < 0) ? 0 : SRC2[111:96];$   
 $DEST[127:112] \leftarrow (SRC2[127:96] > FFFFH) ? FFFFH : TMP[127:112];$   
 $DEST[MAXVL-1:128] \leftarrow 0;$

**VPACKUSDW (VEX.256 encoded version)**

$TMP[15:0] \leftarrow (SRC1[31:0] < 0) ? 0 : SRC1[15:0];$   
 $DEST[15:0] \leftarrow (SRC1[31:0] > FFFFH) ? FFFFH : TMP[15:0];$   
 $TMP[31:16] \leftarrow (SRC1[63:32] < 0) ? 0 : SRC1[47:32];$   
 $DEST[31:16] \leftarrow (SRC1[63:32] > FFFFH) ? FFFFH : TMP[31:16];$   
 $TMP[47:32] \leftarrow (SRC1[95:64] < 0) ? 0 : SRC1[79:64];$   
 $DEST[47:32] \leftarrow (SRC1[95:64] > FFFFH) ? FFFFH : TMP[47:32];$   
 $TMP[63:48] \leftarrow (SRC1[127:96] < 0) ? 0 : SRC1[111:96];$   
 $DEST[63:48] \leftarrow (SRC1[127:96] > FFFFH) ? FFFFH : TMP[63:48];$   
 $TMP[79:64] \leftarrow (SRC2[31:0] < 0) ? 0 : SRC2[15:0];$   
 $DEST[79:64] \leftarrow (SRC2[31:0] > FFFFH) ? FFFFH : TMP[79:64];$   
 $TMP[95:80] \leftarrow (SRC2[63:32] < 0) ? 0 : SRC2[47:32];$   
 $DEST[95:80] \leftarrow (SRC2[63:32] > FFFFH) ? FFFFH : TMP[95:80];$   
 $TMP[111:96] \leftarrow (SRC2[95:64] < 0) ? 0 : SRC2[79:64];$   
 $DEST[111:96] \leftarrow (SRC2[95:64] > FFFFH) ? FFFFH : TMP[111:96];$

```

TMP[127:112] ← (SRC2[127:96] < 0) ? 0 : SRC2[111:96];
DEST[127:112] ← (SRC2[127:96] > FFFFH) ? FFFFH : TMP[127:112];
TMP[143:128] ← (SRC1[159:128] < 0) ? 0 : SRC1[143:128];
DEST[143:128] ← (SRC1[159:128] > FFFFH) ? FFFFH : TMP[143:128];
TMP[159:144] ← (SRC1[191:160] < 0) ? 0 : SRC1[175:160];
DEST[159:144] ← (SRC1[191:160] > FFFFH) ? FFFFH : TMP[159:144];
TMP[175:160] ← (SRC1[223:192] < 0) ? 0 : SRC1[207:192];
DEST[175:160] ← (SRC1[223:192] > FFFFH) ? FFFFH : TMP[175:160];
TMP[191:176] ← (SRC1[255:224] < 0) ? 0 : SRC1[239:224];
DEST[191:176] ← (SRC1[255:224] > FFFFH) ? FFFFH : TMP[191:176];
TMP[207:192] ← (SRC2[159:128] < 0) ? 0 : SRC2[143:128];
DEST[207:192] ← (SRC2[159:128] > FFFFH) ? FFFFH : TMP[207:192];
TMP[223:208] ← (SRC2[191:160] < 0) ? 0 : SRC2[175:160];
DEST[223:208] ← (SRC2[191:160] > FFFFH) ? FFFFH : TMP[223:208];
TMP[239:224] ← (SRC2[223:192] < 0) ? 0 : SRC2[207:192];
DEST[239:224] ← (SRC2[223:192] > FFFFH) ? FFFFH : TMP[239:224];
TMP[255:240] ← (SRC2[255:224] < 0) ? 0 : SRC2[239:224];
DEST[255:240] ← (SRC2[255:224] > FFFFH) ? FFFFH : TMP[255:240];
DEST[MAXVL-1:256] ← 0;

```

**VPACKUSDW (EVEX encoded versions)**

(KL, VL) = (8, 128), (16, 256), (32, 512)

FOR j ← 0 TO ((KL/2) - 1)

  i ← j \* 32

  IF (EVEX.b == 1) AND (SRC2 \*is memory\*)

    THEN

      TMP\_SRC2[j+31:i] ← SRC2[31:0]

    ELSE

      TMP\_SRC2[j+31:i] ← SRC2[j+31:i]

  FI;

ENDFOR;

```

TMP[15:0] ← (SRC1[31:0] < 0) ? 0 : SRC1[15:0];
DEST[15:0] ← (SRC1[31:0] > FFFFH) ? FFFFH : TMP[15:0];
TMP[31:16] ← (SRC1[63:32] < 0) ? 0 : SRC1[47:32];
DEST[31:16] ← (SRC1[63:32] > FFFFH) ? FFFFH : TMP[31:16];
TMP[47:32] ← (SRC1[95:64] < 0) ? 0 : SRC1[79:64];
DEST[47:32] ← (SRC1[95:64] > FFFFH) ? FFFFH : TMP[47:32];
TMP[63:48] ← (SRC1[127:96] < 0) ? 0 : SRC1[111:96];
DEST[63:48] ← (SRC1[127:96] > FFFFH) ? FFFFH : TMP[63:48];
TMP[79:64] ← (TMP_SRC2[31:0] < 0) ? 0 : TMP_SRC2[15:0];
DEST[79:64] ← (TMP_SRC2[31:0] > FFFFH) ? FFFFH : TMP[79:64];
TMP[95:80] ← (TMP_SRC2[63:32] < 0) ? 0 : TMP_SRC2[47:32];
DEST[95:80] ← (TMP_SRC2[63:32] > FFFFH) ? FFFFH : TMP[95:80];
TMP[111:96] ← (TMP_SRC2[95:64] < 0) ? 0 : TMP_SRC2[79:64];
DEST[111:96] ← (TMP_SRC2[95:64] > FFFFH) ? FFFFH : TMP[111:96];
TMP[127:112] ← (TMP_SRC2[127:96] < 0) ? 0 : TMP_SRC2[111:96];
DEST[127:112] ← (TMP_SRC2[127:96] > FFFFH) ? FFFFH : TMP[127:112];
IF VL >= 256
  TMP[143:128] ← (SRC1[159:128] < 0) ? 0 : SRC1[143:128];
  DEST[143:128] ← (SRC1[159:128] > FFFFH) ? FFFFH : TMP[143:128];
  TMP[159:144] ← (SRC1[191:160] < 0) ? 0 : SRC1[175:160];
  DEST[159:144] ← (SRC1[191:160] > FFFFH) ? FFFFH : TMP[159:144];

```

```

TMP[175:160] ← (SRC1[223:192] < 0) ? 0 : SRC1[207:192];
DEST[175:160] ← (SRC1[223:192] > FFFFH) ? FFFFH : TMP[175:160];
TMP[191:176] ← (SRC1[255:224] < 0) ? 0 : SRC1[239:224];
DEST[191:176] ← (SRC1[255:224] > FFFFH) ? FFFFH : TMP[191:176];
TMP[207:192] ← (TMP_SRC2[159:128] < 0) ? 0 : TMP_SRC2[143:128];
DEST[207:192] ← (TMP_SRC2[159:128] > FFFFH) ? FFFFH : TMP[207:192];
TMP[223:208] ← (TMP_SRC2[191:160] < 0) ? 0 : TMP_SRC2[175:160];
DEST[223:208] ← (TMP_SRC2[191:160] > FFFFH) ? FFFFH : TMP[223:208];
TMP[239:224] ← (TMP_SRC2[223:192] < 0) ? 0 : TMP_SRC2[207:192];
DEST[239:224] ← (TMP_SRC2[223:192] > FFFFH) ? FFFFH : TMP[239:224];
TMP[255:240] ← (TMP_SRC2[255:224] < 0) ? 0 : TMP_SRC2[239:224];
DEST[255:240] ← (TMP_SRC2[255:224] > FFFFH) ? FFFFH : TMP[255:240];

```

FI;

IF VL >= 512

```

TMP[271:256] ← (SRC1[287:256] < 0) ? 0 : SRC1[271:256];
DEST[271:256] ← (SRC1[287:256] > FFFFH) ? FFFFH : TMP[271:256];
TMP[287:272] ← (SRC1[319:288] < 0) ? 0 : SRC1[303:288];
DEST[287:272] ← (SRC1[319:288] > FFFFH) ? FFFFH : TMP[287:272];
TMP[303:288] ← (SRC1[351:320] < 0) ? 0 : SRC1[335:320];
DEST[303:288] ← (SRC1[351:320] > FFFFH) ? FFFFH : TMP[303:288];
TMP[319:304] ← (SRC1[383:352] < 0) ? 0 : SRC1[367:352];
DEST[319:304] ← (SRC1[383:352] > FFFFH) ? FFFFH : TMP[319:304];
TMP[335:320] ← (TMP_SRC2[287:256] < 0) ? 0 : TMP_SRC2[271:256];
DEST[335:304] ← (TMP_SRC2[287:256] > FFFFH) ? FFFFH : TMP[79:64];
TMP[351:336] ← (TMP_SRC2[319:288] < 0) ? 0 : TMP_SRC2[303:288];
DEST[351:336] ← (TMP_SRC2[319:288] > FFFFH) ? FFFFH : TMP[351:336];
TMP[367:352] ← (TMP_SRC2[351:320] < 0) ? 0 : TMP_SRC2[315:320];
DEST[367:352] ← (TMP_SRC2[351:320] > FFFFH) ? FFFFH : TMP[367:352];
TMP[383:368] ← (TMP_SRC2[383:352] < 0) ? 0 : TMP_SRC2[367:352];
DEST[383:368] ← (TMP_SRC2[383:352] > FFFFH) ? FFFFH : TMP[383:368];
TMP[399:384] ← (SRC1[415:384] < 0) ? 0 : SRC1[399:384];
DEST[399:384] ← (SRC1[415:384] > FFFFH) ? FFFFH : TMP[399:384];
TMP[415:400] ← (SRC1[447:416] < 0) ? 0 : SRC1[431:416];
DEST[415:400] ← (SRC1[447:416] > FFFFH) ? FFFFH : TMP[415:400];
TMP[431:416] ← (SRC1[479:448] < 0) ? 0 : SRC1[463:448];
DEST[431:416] ← (SRC1[479:448] > FFFFH) ? FFFFH : TMP[431:416];
TMP[447:432] ← (SRC1[511:480] < 0) ? 0 : SRC1[495:480];
DEST[447:432] ← (SRC1[511:480] > FFFFH) ? FFFFH : TMP[447:432];
TMP[463:448] ← (TMP_SRC2[415:384] < 0) ? 0 : TMP_SRC2[399:384];
DEST[463:448] ← (TMP_SRC2[415:384] > FFFFH) ? FFFFH : TMP[463:448];
TMP[475:464] ← (TMP_SRC2[447:416] < 0) ? 0 : TMP_SRC2[431:416];
DEST[475:464] ← (TMP_SRC2[447:416] > FFFFH) ? FFFFH : TMP[475:464];
TMP[491:476] ← (TMP_SRC2[479:448] < 0) ? 0 : TMP_SRC2[463:448];
DEST[491:476] ← (TMP_SRC2[479:448] > FFFFH) ? FFFFH : TMP[491:476];
TMP[511:492] ← (TMP_SRC2[511:480] < 0) ? 0 : TMP_SRC2[495:480];
DEST[511:492] ← (TMP_SRC2[511:480] > FFFFH) ? FFFFH : TMP[511:492];

```

FI;

FOR j ← 0 TO KL-1

  i ← j \* 16

  IF k1[j] OR \*no writemask\*

    THEN

      DEST[i+15:i] ← TMP\_DEST[i+15:i]

    ELSE

      IF \*merging-masking\* ; merging-masking

```

        THEN *DEST[j+15:i] remains unchanged*
        ELSE *zeroing-masking*           ; zeroing-masking
          DEST[j+15:i] ← 0
      FI
  FI;
ENDFOR;
DEST[MAXVL-1:VL] ← 0

```

### Intel C/C++ Compiler Intrinsic Equivalents

```

VPACKUSDW__m512i__mm512_packus_epi32(__m512i m1, __m512i m2);
VPACKUSDW__m512i__mm512_mask_packus_epi32(__m512i s, __mmask32 k, __m512i m1, __m512i m2);
VPACKUSDW__m512i__mm512_maskz_packus_epi32(__mmask32 k, __m512i m1, __m512i m2);
VPACKUSDW__m256i__mm256_mask_packus_epi32(__m256i s, __mmask16 k, __m256i m1, __m256i m2);
VPACKUSDW__m256i__mm256_maskz_packus_epi32(__mmask16 k, __m256i m1, __m256i m2);
VPACKUSDW__m128i__mm_mask_packus_epi32(__m128i s, __mmask8 k, __m128i m1, __m128i m2);
VPACKUSDW__m128i__mm_maskz_packus_epi32(__mmask8 k, __m128i m1, __m128i m2);
PACKUSDW__m128i__mm_packus_epi32(__m128i m1, __m128i m2);
VPACKUSDW__m256i__mm256_packus_epi32(__m256i m1, __m256i m2);

```

### SIMD Floating-Point Exceptions

None

### Other Exceptions

Non-EVEX-encoded instruction, see Exceptions Type 4.

EVEX-encoded instruction, see Exceptions Type E4NF.

## PTWRITE - Write Data to a Processor Trace Packet

| Opcode/<br>Instruction               | Op/<br>En | 64/32 bit<br>Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|--------------------------------------|-----------|------------------------------|--------------------------|--|
| F3 REX.W OF AE /4<br>PTWRITE r64/m64 | RM        | V/N.E                        |                          | Reads the data from r64/m64 to encode into a PTW packet if dependencies are met (see details below). |
| F3 OF AE /4<br>PTWRITE r32/m32       | RM        | V/V                          |                          | Reads the data from r32/m32 to encode into a PTW packet if dependencies are met (see details below). |

### Instruction Operand Encoding

| Op/En | Operand 1    | Operand 2 | Operand 3 | Operand 4 |
|-------|--------------|-----------|-----------|-----------|
| RM    | ModRM:rm (r) | NA        | NA        | NA        |

### Description

This instruction reads data in the source operand and sends it to the Intel Processor Trace hardware to be encoded in a PTW packet if TriggerEn, ContextEn, FilterEn, and PTWEn are all set to 1. For more details on these values, see *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C, Section 35.2.2, "Software Trace Instrumentation with PTWRITE"*. The size of data is 64-bit if using REX.W in 64-bit mode, otherwise 32-bits of data are copied from the source operand.

Note: The instruction will #UD if prefix 66H is used.

### Operation

IF (IA32\_RTIT\_STATUS.TriggerEn & IA32\_RTIT\_STATUS.ContextEn & IA32\_RTIT\_STATUS.FilterEn & IA32\_RTIT\_CTL.PTWEn) = 1

PTW.PayloadBytes ← Encoded payload size;

PTW.IP ← IA32\_RTIT\_CTL.FUPonPTW

IF IA32\_RTIT\_CTL.FUPonPTW = 1

Insert FUP packet with IP of PTWRITE;

FI;

FI;

### Flags Affected

None.

### Other Exceptions

- #GP(0) If a memory operand effective address is outside the CS, DS, ES, FS or GS segments.
- #SS(0) If a memory operand effective address is outside the SS segment limit.
- #PF (fault-code) For a page fault.
- #AC(0) If an unaligned memory reference is made while the current privilege level is 3 and alignment checking is enabled.
- #UD If CPUID.(EAX=14H, ECX=0):EBX.PTWRITE [Bit 4] = 0.  
If LOCK prefix is used.  
If 66H prefix is used.

**Real-Address Mode Exceptions**

|        |   |
|--------|---|
| #GP(0) | If any part of the operand lies outside of the effective address space from 0 to 0FFFFH.                |
| #SS(0) | If a memory operand effective address is outside the SS segment limit.                                  |
| #UD    | If CPUID.(EAX=14H, ECX=0):EBX.PTWRITE [Bit 4] = 0.<br>If LOCK prefix is used.<br>If 66H prefix is used. |

**Virtual 8086 Mode Exceptions**

|                  |   |
|------------------|---|
| #GP(0)           | If any part of the operand lies outside of the effective address space from 0 to 0FFFFH.                |
| #SS(0)           | If a memory operand effective address is outside the SS segment limit.                                  |
| #PF (fault-code) | For a page fault.   |
| #AC(0)           | If an unaligned memory reference is made while alignment checking is enabled.                           |
| #UD              | If CPUID.(EAX=14H, ECX=0):EBX.PTWRITE [Bit 4] = 0.<br>If LOCK prefix is used.<br>If 66H prefix is used. |

**Compatibility Mode Exceptions**

Same exceptions as in Protected Mode.

**64-Bit Mode Exceptions**

|                  |  |
|------------------|--|
| #GP(0)           | If the memory address is in a non-canonical form.  |
| #SS(0)           | If a memory address referencing the SS segment is in a non-canonical form.   |
| #PF (fault-code) | For a page fault.  |
| #AC(0)           | If alignment checking is enabled and an unaligned memory reference is made while the current privilege level is 3. |
| #UD              | If CPUID.(EAX=14H, ECX=0):EBX.PTWRITE [Bit 4] = 0.<br>If LOCK prefix is used.<br>If 66H prefix is used.            |

## SQRTPD—Square Root of Double-Precision Floating-Point Values

| Opcode/<br>Instruction  | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|---|-----------|------------------------------|--------------------------|--|
| 66 0F 51 /r<br>SQRTPD xmm1, xmm2/m128                                     | A         | V/V                          | SSE2                     | Computes Square Roots of the packed double-precision floating-point values in xmm2/m128 and stores the result in xmm1.                                 |
| VEX.128.66.0F.WIG 51 /r<br>VSQRTPD xmm1, xmm2/m128                        | A         | V/V                          | AVX                      | Computes Square Roots of the packed double-precision floating-point values in xmm2/m128 and stores the result in xmm1.                                 |
| VEX.256.66.0F.WIG 51 /r<br>VSQRTPD ymm1, ymm2/m256                        | A         | V/V                          | AVX                      | Computes Square Roots of the packed double-precision floating-point values in ymm2/m256 and stores the result in ymm1.                                 |
| EVEX.128.66.0F.W1 51 /r<br>VSQRTPD xmm1 {k1}{z},<br>xmm2/m128/m64bcst     | B         | V/V                          | AVX512VL<br>AVX512F      | Computes Square Roots of the packed double-precision floating-point values in xmm2/m128/m64bcst and stores the result in xmm1 subject to writemask k1. |
| EVEX.256.66.0F.W1 51 /r<br>VSQRTPD ymm1 {k1}{z},<br>ymm2/m256/m64bcst     | B         | V/V                          | AVX512VL<br>AVX512F      | Computes Square Roots of the packed double-precision floating-point values in ymm2/m256/m64bcst and stores the result in ymm1 subject to writemask k1. |
| EVEX.512.66.0F.W1 51 /r<br>VSQRTPD zmm1 {k1}{z},<br>zmm2/m512/m64bcst{er} | B         | V/V                          | AVX512F                  | Computes Square Roots of the packed double-precision floating-point values in zmm2/m512/m64bcst and stores the result in zmm1 subject to writemask k1. |

### Instruction Operand Encoding

| Op/En | Tuple Type | Operand 1     | Operand 2     | Operand 3 | Operand 4 |
|-------|------------|---------------|---------------|-----------|-----------|
| A     | NA         | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |
| B     | Full       | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |

### Description

Performs a SIMD computation of the square roots of the two, four or eight packed double-precision floating-point values in the source operand (the second operand) stores the packed double-precision floating-point results in the destination operand (the first operand).

EVEX encoded versions: The source operand is a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 64-bit memory location. The destination operand is a ZMM/YMM/XMM register updated according to the writemask.

VEX.256 encoded version: The source operand is a YMM register or a 256-bit memory location. The destination operand is a YMM register. The upper bits (MAXVL-1:256) of the corresponding ZMM register destination are zeroed.

VEX.128 encoded version: the source operand second source operand or a 128-bit memory location. The destination operand is an XMM register. The upper bits (MAXVL-1:128) of the corresponding ZMM register destination are zeroed.

128-bit Legacy SSE version: The second source can be an XMM register or 128-bit memory location. The destination is not distinct from the first source XMM register and the upper bits (MAXVL-1:128) of the corresponding ZMM register destination are unmodified.

Note: VEX.vvvv and EVEX.vvvv are reserved and must be 1111b otherwise instructions will #UD.



**Operation****VSQRTPD (EVEX encoded versions)**

(KL, VL) = (2, 128), (4, 256), (8, 512)

IF (VL = 512) AND (EVEX.b = 1) AND (SRC \*is register\*)

THEN

SET\_RM(EVEX.RC);

ELSE

SET\_RM(MXCSR.RM);

FI;

FOR j ← 0 TO KL-1

i ← j \* 64

IF k1[j] OR \*no writemask\* THEN

IF (EVEX.b = 1) AND (SRC \*is memory\*)

THEN DEST[i+63:i] ← SQRT(SRC[63:0])

ELSE DEST[i+63:i] ← SQRT(SRC[i+63:i])

FI;

ELSE

IF \*merging-masking\* ; merging-masking

THEN \*DEST[i+63:i] remains unchanged\*

ELSE ; zeroing-masking

DEST[i+63:i] ← 0

FI

FI;

ENDFOR

DEST[MAXVL-1:VL] ← 0

**VSQRTPD (VEX.256 encoded version)**

DEST[63:0] ← SQRT(SRC[63:0])

DEST[127:64] ← SQRT(SRC[127:64])

DEST[191:128] ← SQRT(SRC[191:128])

DEST[255:192] ← SQRT(SRC[255:192])

DEST[MAXVL-1:256] ← 0

**VSQRTPD (VEX.128 encoded version)**

DEST[63:0] ← SQRT(SRC[63:0])

DEST[127:64] ← SQRT(SRC[127:64])

DEST[MAXVL-1:128] ← 0

**SQRTPD (128-bit Legacy SSE version)**

DEST[63:0] ← SQRT(SRC[63:0])

DEST[127:64] ← SQRT(SRC[127:64])

DEST[MAXVL-1:128] (Unmodified)

**Intel C/C++ Compiler Intrinsic Equivalent**

VSQRTPD \_\_m512d \_\_mm512\_sqrt\_round\_pd(\_\_m512d a, int r);

VSQRTPD \_\_m512d \_\_mm512\_mask\_sqrt\_round\_pd(\_\_m512d s, \_\_mmask8 k, \_\_m512d a, int r);

VSQRTPD \_\_m512d \_\_mm512\_maskz\_sqrt\_round\_pd(\_\_mmask8 k, \_\_m512d a, int r);

VSQRTPD \_\_m256d \_\_mm256\_sqrt\_pd(\_\_m256d a);

VSQRTPD \_\_m256d \_\_mm256\_mask\_sqrt\_pd(\_\_m256d s, \_\_mmask8 k, \_\_m256d a, int r);

VSQRTPD \_\_m256d \_\_mm256\_maskz\_sqrt\_pd(\_\_mmask8 k, \_\_m256d a, int r);

SQRTPD \_\_m128d \_\_mm\_sqrt\_pd(\_\_m128d a);

VSQRTPD \_\_m128d \_\_mm\_mask\_sqrt\_pd(\_\_m128d s, \_\_mmask8 k, \_\_m128d a, int r);

VSQRTPD \_\_m128d \_\_mm\_maskz\_sqrt\_pd(\_\_mmask8 k, \_\_m128d a, int r);

**SIMD Floating-Point Exceptions**

Invalid, Precision, Denormal

**Other Exceptions**

Non-EVEX-encoded instruction, see Exceptions Type 2; additionally

#UD                    If VEX.vvvv != 1111B.

EVEX-encoded instruction, see Exceptions Type E2.

#UD                    If EVEX.vvvv != 1111B.

## 6. Updates to Chapter 5, Volume 2C

Change bars show changes to Chapter 5 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C: Instruction Set Reference, V-Z*.

-----  
Change to this chapter: Updates to VGETEXPPD, VGETEXPPS and WBINVD instructions.

## VGETEXPPD—Convert Exponents of Packed DP FP Values to DP FP Values

| Opcode/<br>Instruction   | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|--|-----------|------------------------------|--------------------------|--|
| EVEX.128.66.0F38.W1 42 /r<br>VGETEXPPD xmm1 {k1}{z},<br>xmm2/m128/m64bcst      | A         | V/V                          | AVX512VL<br>AVX512F      | Convert the exponent of packed double-precision floating-point values in the source operand to DP FP results representing unbiased integer exponents and stores the results in the destination register.           |
| EVEX.256.66.0F38.W1 42 /r<br>VGETEXPPD ymm1 {k1}{z},<br>ymm2/m256/m64bcst      | A         | V/V                          | AVX512VL<br>AVX512F      | Convert the exponent of packed double-precision floating-point values in the source operand to DP FP results representing unbiased integer exponents and stores the results in the destination register.           |
| EVEX.512.66.0F38.W1 42 /r<br>VGETEXPPD zmm1 {k1}{z},<br>zmm2/m512/m64bcst{sae} | A         | V/V                          | AVX512F                  | Convert the exponent of packed double-precision floating-point values in the source operand to DP FP results representing unbiased integer exponents and stores the results in the destination under writemask k1. |

### Instruction Operand Encoding

| Op/En | Tuple Type | Operand 1     | Operand 2     | Operand 3 | Operand 4 |
|-------|------------|---------------|---------------|-----------|-----------|
| A     | Full       | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |

### Description

Extracts the biased exponents from the normalized DP FP representation of each qword data element of the source operand (the second operand) as unbiased signed integer value, or convert the denormal representation of input data to unbiased negative integer values. Each integer value of the unbiased exponent is converted to double-precision FP value and written to the corresponding qword elements of the destination operand (the first operand) as DP FP numbers.

The destination operand is a ZMM/YMM/XMM register and updated under the writemask. The source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 64-bit memory location.

EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

Each GETEXP operation converts the exponent value into a FP number (permitting input value in denormal representation). Special cases of input values are listed in Table 5-5.

The formula is:

$$\text{GETEXP}(x) = \text{floor}(\log_2(|x|))$$

Notation **floor(x)** stands for the greatest integer not exceeding real number x.

**Table 5-5. VGETEXPPD/SD Special Cases**

| Input Operand                    | Result                                | Comments   |
|----------------------------------|---------------------------------------|--|
| src1 = NaN                       | QNaN(src1)                            | If (SRC = SNaN) then #IE<br>If (SRC = denormal) then #DE |
| $0 <  \text{src1}  < \text{INF}$ | $\text{floor}(\log_2( \text{src1} ))$ |  |
| $ \text{src1}  = +\text{INF}$    | +INF                                  |  |
| $ \text{src1}  = 0$              | -INF                                  |  |

**Operation**

NormalizeExpTinyDPFP(SRC[63:0])

```
{
  // Jbit is the hidden integral bit of a FP number. In case of denormal number it has the value of ZERO.
  Src.Jbit ← 0;
  Dst.exp ← 1;
  Dst.fraction ← SRC[51:0];
  WHILE(Src.Jbit = 0)
  {
    Src.Jbit ← Dst.fraction[51];      // Get the fraction MSB
    Dst.fraction ← Dst.fraction << 1; // One bit shift left
    Dst.exp--;                       // Decrement the exponent
  }
  Dst.fraction ← 0;                  // zero out fraction bits
  Dst.sign ← 1;                      // Return negative sign
  TMP[63:0] ← MXCSR.DAZ? 0 : (Dst.sign << 63) OR (Dst.exp << 52) OR (Dst.fraction);
  Return (TMP[63:0]);
}
```

ConvertExpDPFP(SRC[63:0])

```
{
  Src.sign ← 0;                      // Zero out sign bit
  Src.exp ← SRC[62:52];
  Src.fraction ← SRC[51:0];
  // Check for NaN
  IF (SRC = NaN)
  {
    IF ( SRC = SNAN ) SET IE;
    Return QNAN(SRC);
  }
  // Check for +INF
  IF (SRC = +INF) Return (SRC);

  // check if zero operand
  IF ((Src.exp = 0) AND ((Src.fraction = 0) OR (MXCSR.DAZ = 1))) Return (-INF);
}
ELSE // check if denormal operand (notice that MXCSR.DAZ = 0)
{
  IF ((Src.exp = 0) AND (Src.fraction != 0))
  {
    TMP[63:0] ← NormalizeExpTinyDPFP(SRC[63:0]); // Get Normalized Exponent
    Set #DE
  }
  ELSE // exponent value is correct
  {
    TMP[63:0] ← (Src.sign << 63) OR (Src.exp << 52) OR (Src.fraction);
  }
  TMP ← SAR(TMP, 52); // Shift Arithmetic Right
  TMP ← TMP - 1023; // Subtract Bias
  Return CvtI2D(TMP); // Convert INT to Double-Precision FP number
}
}
```

**VGETEXPPD (EVEX encoded versions)**

(KL, VL) = (2, 128), (4, 256), (8, 512)

FOR j ← 0 TO KL-1

i ← j \* 64

IF k1[j] OR \*no writemask\*

THEN

IF (EVEX.b = 1) AND (SRC \*is memory\*)

THEN

DEST[i+63:i] ←

ConvertExpDPFP(SRC[63:0])

ELSE

DEST[i+63:i] ←

ConvertExpDPFP(SRC[i+63:i])

FI;

ELSE

IF \*merging-masking\* ; merging-masking

THEN \*DEST[i+63:i] remains unchanged\*

ELSE ; zeroing-masking

DEST[i+63:i] ← 0

FI

FI;

ENDFOR

DEST[MAXVL-1:VL] ← 0

**Intel C/C++ Compiler Intrinsic Equivalent**

VGETEXPPD \_\_m512d \_mm512\_getexp\_pd(\_\_m512d a);

VGETEXPPD \_\_m512d \_mm512\_mask\_getexp\_pd(\_\_m512d s, \_\_mmask8 k, \_\_m512d a);

VGETEXPPD \_\_m512d \_mm512\_maskz\_getexp\_pd(\_\_mmask8 k, \_\_m512d a);

VGETEXPPD \_\_m512d \_mm512\_getexp\_round\_pd(\_\_m512d a, int sae);

VGETEXPPD \_\_m512d \_mm512\_mask\_getexp\_round\_pd(\_\_m512d s, \_\_mmask8 k, \_\_m512d a, int sae);

VGETEXPPD \_\_m512d \_mm512\_maskz\_getexp\_round\_pd(\_\_mmask8 k, \_\_m512d a, int sae);

VGETEXPPD \_\_m256d \_mm256\_getexp\_pd(\_\_m256d a);

VGETEXPPD \_\_m256d \_mm256\_mask\_getexp\_pd(\_\_m256d s, \_\_mmask8 k, \_\_m256d a);

VGETEXPPD \_\_m256d \_mm256\_maskz\_getexp\_pd(\_\_mmask8 k, \_\_m256d a);

VGETEXPPD \_\_m128d \_mm\_getexp\_pd(\_\_m128d a);

VGETEXPPD \_\_m128d \_mm\_mask\_getexp\_pd(\_\_m128d s, \_\_mmask8 k, \_\_m128d a);

VGETEXPPD \_\_m128d \_mm\_maskz\_getexp\_pd(\_\_mmask8 k, \_\_m128d a);

**SIMD Floating-Point Exceptions**

Invalid, Denormal

**Other Exceptions**

See Exceptions Type E2.

#UD If EVEX.vvvv != 1111B.

## VGETEXPPS—Convert Exponents of Packed SP FP Values to SP FP Values

| Opcode/<br>Instruction   | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|--|-----------|------------------------------|--------------------------|--|
| EVEX.128.66.0F38.W0 42 /r<br>VGETEXPPS xmm1 {k1}{z},<br>xmm2/m128/m32bcst      | A         | V/V                          | AVX512VL<br>AVX512F      | Convert the exponent of packed single-precision floating-point values in the source operand to SP FP results representing unbiased integer exponents and stores the results in the destination register. |
| EVEX.256.66.0F38.W0 42 /r<br>VGETEXPPS ymm1 {k1}{z},<br>ymm2/m256/m32bcst      | A         | V/V                          | AVX512VL<br>AVX512F      | Convert the exponent of packed single-precision floating-point values in the source operand to SP FP results representing unbiased integer exponents and stores the results in the destination register. |
| EVEX.512.66.0F38.W0 42 /r<br>VGETEXPPS zmm1 {k1}{z},<br>zmm2/m512/m32bcst{sae} | A         | V/V                          | AVX512F                  | Convert the exponent of packed single-precision floating-point values in the source operand to SP FP results representing unbiased integer exponents and stores the results in the destination register. |

### Instruction Operand Encoding

| Op/En | Tuple Type | Operand 1     | Operand 2     | Operand 3 | Operand 4 |
|-------|------------|---------------|---------------|-----------|-----------|
| A     | Full       | ModRM:reg (w) | ModRM:r/m (r) | NA        | NA        |

### Description

Extracts the biased exponents from the normalized SP FP representation of each dword element of the source operand (the second operand) as unbiased signed integer value, or convert the denormal representation of input data to unbiased negative integer values. Each integer value of the unbiased exponent is converted to single-precision FP value and written to the corresponding dword elements of the destination operand (the first operand) as SP FP numbers.

The destination operand is a ZMM/YMM/XMM register and updated under the writemask. The source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 32-bit memory location.

EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

Each GETEXP operation converts the exponent value into a FP number (permitting input value in denormal representation). Special cases of input values are listed in Table 5-6.

The formula is:

$$\text{GETEXP}(x) = \text{floor}(\log_2(|x|))$$

Notation **floor(x)** stands for maximal integer not exceeding real number x.

Software usage of VGETEXPxx and VGETMANTxx instructions generally involve a combination of GETEXP operation and GETMANT operation (see VGETMANTPD). Thus VGETEXPxx instruction do not require software to handle SIMD FP exceptions.

Table 5-6. VGETEXPPS/SS Special Cases

| Input Operand    | Result                           | Comments   |
|------------------|----------------------------------|--|
| src1 = NaN       | QNaN(src1)                       | If (SRC = SNaN) then #IE<br>If (SRC = denormal) then #DE |
| 0 <  src1  < INF | floor(log <sub>2</sub> ( src1 )) |  |
| src1  = +INF     | +INF                             |  |
| src1  = 0        | -INF                             |  |

Figure 5-14 illustrates the VGETEXPPS functionality on input values with normalized representation.

|                                 | 31 | 30  | 29 | 28 | 27 | 26 | 25 | 24 | 23 | 22       | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |   |
|---------------------------------|----|-----|----|----|----|----|----|----|----|----------|----|----|----|----|----|----|----|----|----|----|----|----|---|---|---|---|---|---|---|---|---|---|---|
|                                 | s  | exp |    |    |    |    |    |    |    | Fraction |    |    |    |    |    |    |    |    |    |    |    |    |   |   |   |   |   |   |   |   |   |   |   |
| Src = 2 <sup>-1</sup>           | 0  | 1   | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0        | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| SAR Src, 23 = 080h              | 0  | 0   | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0        | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| -Bias                           | 1  | 1   | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1        | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1  | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 1 |   |
| Tmp - Bias = 1                  | 0  | 0   | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0        | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 |
| Cvt_P12PS(01h) = 2 <sup>0</sup> | 0  | 0   | 1  | 1  | 1  | 1  | 1  | 1  | 0  | 0        | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Figure 5-14. VGETEXPPS Functionality On Normal Input values

### Operation

NormalizeExpTinySPFP(SRC[31:0])

```
{
    // Jbit is the hidden integral bit of a FP number. In case of denormal number it has the value of ZERO.
    Src.Jbit ← 0;
    Dst.exp ← 1;
    Dst.fraction ← SRC[22:0];
    WHILE(Src.Jbit = 0)
    {
        Src.Jbit ← Dst.fraction[22]; // Get the fraction MSB
        Dst.fraction ← Dst.fraction << 1; // One bit shift left
        Dst.exp--; // Decrement the exponent
    }
    Dst.fraction ← 0; // zero out fraction bits
    Dst.sign ← 1; // Return negative sign
    TMP[31:0] ← MXCSR.DAZ? 0 : (Dst.sign << 31) OR (Dst.exp << 23) OR (Dst.fraction);
    Return (TMP[31:0]);
}
```

ConvertExpSPFP(SRC[31:0])

```
{
    Src.sign ← 0; // Zero out sign bit
    Src.exp ← SRC[30:23];
    Src.fraction ← SRC[22:0];
    // Check for NaN
    IF (SRC = NaN)
    {
        IF (SRC = SNAN) SET IE;
        Return QNAN(SRC);
    }
    // Check for +INF
    IF (SRC = +INF) Return (SRC);

    // check if zero operand
    IF ((Src.exp = 0) AND ((Src.fraction = 0) OR (MXCSR.DAZ = 1))) Return (-INF);
}
ELSE // check if denormal operand (notice that MXCSR.DAZ = 0)
{
```



```

    IF ((Src.exp = 0) AND (Src.fraction != 0))
    {
        TMP[31:0] ← NormalizeExpTinySPFP(SRC[31:0]);    // Get Normalized Exponent
        Set #DE
    }
    ELSE // exponent value is correct
    {
        TMP[31:0] ← (Src.sign << 31) OR (Src.exp << 23) OR (Src.fraction);
    }
    TMP ← SAR(TMP, 23); // Shift Arithmetic Right
    TMP ← TMP - 127; // Subtract Bias
    Return CvtI2D(TMP); // Convert INT to Single-Precision FP number
}
}

```

**VGETEXPPS (EVEX encoded versions)**

(KL, VL) = (4, 128), (8, 256), (16, 512)

FOR j ← 0 TO KL-1

i ← j \* 32

IF k1[j] OR \*no writemask\*

THEN

IF (EVEX.b = 1) AND (SRC \*is memory\*)

THEN

DEST[i+31:i] ←

ConvertExpSPFP(SRC[31:0])

ELSE

DEST[i+31:i] ←

ConvertExpSPFP(SRC[i+31:i])

FI;

ELSE

IF \*merging-masking\* ; merging-masking

THEN \*DEST[i+31:i] remains unchanged\*

ELSE ; zeroing-masking

DEST[i+31:i] ← 0

FI

FI;

ENDFOR

DEST[MAXVL-1:VL] ← 0

**Intel C/C++ Compiler Intrinsic Equivalent**

```

VGETEXPPS __m512 _mm512_getexp_ps(__m512 a);
VGETEXPPS __m512 _mm512_mask_getexp_ps(__m512 s, __mmask16 k, __m512 a);
VGETEXPPS __m512 _mm512_maskz_getexp_ps(__mmask16 k, __m512 a);
VGETEXPPS __m512 _mm512_getexp_round_ps(__m512 a, int sae);
VGETEXPPS __m512 _mm512_mask_getexp_round_ps(__m512 s, __mmask16 k, __m512 a, int sae);
VGETEXPPS __m512 _mm512_maskz_getexp_round_ps(__mmask16 k, __m512 a, int sae);
VGETEXPPS __m256 _mm256_getexp_ps(__m256 a);
VGETEXPPS __m256 _mm256_mask_getexp_ps(__m256 s, __mmask8 k, __m256 a);
VGETEXPPS __m256 _mm256_maskz_getexp_ps(__mmask8 k, __m256 a);
VGETEXPPS __m128 _mm_getexp_ps(__m128 a);
VGETEXPPS __m128 _mm_mask_getexp_ps(__m128 s, __mmask8 k, __m128 a);
VGETEXPPS __m128 _mm_maskz_getexp_ps(__mmask8 k, __m128 a);

```

**SIMD Floating-Point Exceptions**

Invalid, Denormal

**Other Exceptions**

See Exceptions Type E2.

#UD                    If EVEX.vvvv != 1111B.

## WBINVD—Write Back and Invalidate Cache

| Opcode | Instruction | Op/En | 64-Bit Mode | Compat/Leg Mode | Description  |
|--------|-------------|-------|-------------|-----------------|--|
| 0F 09  | WBINVD      | Z0    | Valid       | Valid           | Write back and flush Internal caches; initiate writing-back and flushing of external caches. |

### Instruction Operand Encoding

| Op/En | Operand 1 | Operand 2 | Operand 3 | Operand 4 |
|-------|-----------|-----------|-----------|-----------|
| Z0    | NA        | NA        | NA        | NA        |

### Description

Writes back all modified cache lines in the processor's internal cache to main memory and invalidates (flushes) the internal caches. The instruction then issues a special-function bus cycle that directs external caches to also write back modified data and another bus cycle to indicate that the external caches should be invalidated.

After executing this instruction, the processor does not wait for the external caches to complete their write-back and flushing operations before proceeding with instruction execution. It is the responsibility of hardware to respond to the cache write-back and flush signals. The amount of time or cycles for WBINVD to complete will vary due to size and other factors of different cache hierarchies. As a consequence, the use of the WBINVD instruction can have an impact on logical processor interrupt/event response time. Additional information of WBINVD behavior in a cache hierarchy with hierarchical sharing topology can be found in Chapter 2 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.

The WBINVD instruction is a privileged instruction. When the processor is running in protected mode, the CPL of a program or procedure must be 0 to execute this instruction. This instruction is also a serializing instruction (see "Serializing Instructions" in Chapter 8 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*).

In situations where cache coherency with main memory is not a concern, software can use the INVD instruction. This instruction's operation is the same in non-64-bit modes and 64-bit mode.

### IA-32 Architecture Compatibility

The WBINVD instruction is implementation dependent, and its function may be implemented differently on future Intel 64 and IA-32 processors. The instruction is not supported on IA-32 processors earlier than the Intel486 processor.

### Operation

```
WriteBack(InternalCaches);
Flush(InternalCaches);
SignalWriteBack(ExternalCaches);
SignalFlush(ExternalCaches);
Continue; (* Continue execution *)
```

### Intel C/C++ Compiler Intrinsic Equivalent

```
WBINVD void _wbinvd(void);
```

### Flags Affected

None.

### Protected Mode Exceptions

```
#GP(0)      If the current privilege level is not 0.
#UD        If the LOCK prefix is used.
```

**Real-Address Mode Exceptions**

#UD If the LOCK prefix is used.

**Virtual-8086 Mode Exceptions**

#GP(0) WBINVD cannot be executed at the virtual-8086 mode.

**Compatibility Mode Exceptions**

Same exceptions as in protected mode.

**64-Bit Mode Exceptions**

Same exceptions as in protected mode.

## 7. Updates to Chapter 7, Volume 2D

Change bars show changes to Chapter 7 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2D: Instruction Set Reference*.

-----  
Change to this chapter: Added V4FMADDPS/V4FNMADDPS, V4FMADDSS/V4FNMADDSS and VP4DPWSSD instructions.

# CHAPTER 7 INSTRUCTION SET REFERENCE UNIQUE TO INTEL® XEON PHI™ PROCESSORS

---

This chapter describes the instruction set that is unique to Intel® Xeon Phi™ Processors based on the Knights Landing and Knights Mill microarchitectures. The set is not supported in any other Intel processors. Included are Intel® AVX-512 instructions. For additional instructions supported on these processors, see Chapter 3, “Instruction Set Reference, A-L”, Chapter 4, “Instruction Set Reference, M-U”, and Chapter 5, “Instruction Set Reference, V-Z”.

## V4FMADDPS/V4FNMADDPS — Packed Single-Precision Floating-Point Fused Multiply-Add (4-iterations)

| Opcode/<br>Instruction   | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID Feature<br>Flag | Description   |
|--|-----------|------------------------------|-----------------------|---|
| EVEX.DDS.512.F2.0F38.W0 9A /r<br>V4FMADDPS zmm1{k1}{z}, zmm2+3,<br>m128  | A         | V/V                          | AVX512_4FMAPS         | Multiply packed single-precision floating-point values from source register block indicated by zmm2 by values from m128 and accumulate the result in zmm1.            |
| EVEX.DDS.512.F2.0F38.W0 AA /r<br>V4FNMADDPS zmm1{k1}{z},<br>zmm2+3, m128 | A         | V/V                          | AVX512_4FMAPS         | Multiply and negate packed single-precision floating-point values from source register block indicated by zmm2 by values from m128 and accumulate the result in zmm1. |

### Instruction Operand Encoding

| Op/En | Tuple     | Operand 1        | Operand 2     | Operand 3     | Operand 4 |
|-------|-----------|------------------|---------------|---------------|-----------|
| A     | Tuple1_4X | ModRM:reg (r, w) | EVEX.vvvv (r) | ModRM:r/m (r) | NA        |

### Description

This instruction computes 4 sequential packed fused single-precision floating-point multiply-add instructions with a sequentially selected memory operand in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based on that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if any of the 16 lowest significant mask bits is set to 1 or if a “no masking” encoding is used.

The tuple type Tuple1\_4X implies that four 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

Rounding is performed at every FMA (fused multiply and add) boundary. Exceptions are also taken sequentially. Pre- and post-computational exceptions of the first FMA take priority over the pre- and post-computational exceptions of the second FMA, etc.

## Operation

src\_reg\_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

```
define NFMA_PS(kl, vl, dest, k1, msrc, regs_loaded, src_base, posneg):
```

```
    tmpdest ← dest
```

```
    // reg[] is an array representing the SIMD register file.
```

```
    FOR j ← 0 to regs_loaded-1:
```

```
        FOR i ← 0 to kl-1:
```

```
            IF k1[i] or *no writemask*:
```

```
                IF posneg = 0:
```

```
                    tmpdest.single[i] ← RoundFPControl_MXCSR(tmpdest.single[i] - reg[src_base + j].single[i] * msrc.single[j])
```

```
                ELSE:
```

```
                    tmpdest.single[i] ← RoundFPControl_MXCSR(tmpdest.single[i] + reg[src_base + j].single[i] * msrc.single[j])
```

```
            ELSE IF *zeroing*:
```

```
                tmpdest.single[i] ← 0
```

```
    dest ← tmpdst
```

```
    dest[MAX_VL-1:VL] ← 0
```

V4FMADDPS and V4FNMADDPS dest{k1}, src1, msrc (AVX512)

KL, VL = (16,512)

```
regs_loaded ← 4
```

```
src_base ← src_reg_id & ~3 // for src1 operand
```

```
posneg ← 0 if negative form, 1 otherwise
```

```
NFMA_PS(kl, vl, dest, k1, msrc, regs_loaded, src_base, posneg)
```

## Intel C/C++ Compiler Intrinsic Equivalent

```
V4FMADDPS __m512 __mm512_4fmadd_ps(__m512, __m512x4, __m128 *);
```

```
V4FMADDPS __m512 __mm512_mask_4fmadd_ps(__m512, __mmask16, __m512x4, __m128 *);
```

```
V4FMADDPS __m512 __mm512_maskz_4fmadd_ps(__mmask16, __m512, __m512x4, __m128 *);
```

```
V4FNMADDPS __m512 __mm512_4fnmadd_ps(__m512, __m512x4, __m128 *);
```

```
V4FNMADDPS __m512 __mm512_mask_4fnmadd_ps(__m512, __mmask16, __m512x4, __m128 *);
```

```
V4FNMADDPS __m512 __mm512_maskz_4fnmadd_ps(__mmask16, __m512, __m512x4, __m128 *);
```

## SIMD Floating-Point Exceptions

Overflow, Underflow, Invalid, Precision, Denormal.

## Other Exceptions

See Type E2; additionally

#UD If the EVEX broadcast bit is set to 1.

#UD If the MODRM.mod = 0b11.



## V4FMADDSS/V4FNMADDSS —Scalar Single-Precision Floating-Point Fused Multiply-Add (4-iterations)

| Opcode/<br>Instruction  | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID Feature<br>Flag | Description   |
|---|-----------|------------------------------|-----------------------|---|
| EVEX.DDS.LLIG.F2.0F38.W0 9B /r<br>V4FMADDSS xmm1{k1}{z},<br>xmm2+3, m128  | A         | V/V                          | AVX512_4FMAPS         | Multiply scalar single-precision floating-point values from source register block indicated by xmm2 by values from m128 and accumulate the result in xmm1.            |
| EVEX.DDS.LLIG.F2.0F38.W0 AB /r<br>V4FNMADDSS xmm1{k1}{z},<br>xmm2+3, m128 | A         | V/V                          | AVX512_4FMAPS         | Multiply and negate scalar single-precision floating-point values from source register block indicated by xmm2 by values from m128 and accumulate the result in xmm1. |

### Instruction Operand Encoding

| Op/En | Tuple     | Operand 1        | Operand 2     | Operand 3     | Operand 4 |
|-------|-----------|------------------|---------------|---------------|-----------|
| A     | Tuple1_4X | ModRM:reg (r, w) | EVEX.vvvv (r) | ModRM:r/m (r) | NA        |

### Description

This instruction computes 4 sequential scalar fused single-precision floating-point multiply-add instructions with a sequentially selected memory operand in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if the least significant mask bit is set to 1 or if a “no masking” encoding is used.

The tuple type Tuple1\_4X implies that four 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

Rounding is performed at every FMA boundary. Exceptions are also taken sequentially. Pre- and post-computational exceptions of the first FMA take priority over the pre- and post-computational exceptions of the second FMA, etc.

### Operation

src\_reg\_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

```
define NFMA_SS(vl, dest, k1, msrc, regs_loaded, src_base, posneg):
    tmpdest ← dest
    // reg[] is an array representing the SIMD register file.
    IF k1[0] or *no writemask*:
        FOR j ← 0 to regs_loaded - 1:
            IF posneg = 0:
                tmpdest.single[0] ← RoundFPControl_MXCSR(tmpdest.single[0] - reg[src_base + j].single[0] * msrc.single[j])
            ELSE:
                tmpdest.single[0] ← RoundFPControl_MXCSR(tmpdest.single[0] + reg[src_base + j].single[0] * msrc.single[j])
    ELSE IF *zeroing*:
        tmpdest.single[0] ← 0
    dest ← tmpdst
    dest[MAX_VL-1:VL] ← 0
```

V4FMADDSS and V4FNMADDSS dest{k1}, src1, msrc (AVX512)  
 VL = 128

regs\_loaded ← 4  
 src\_base ← src\_reg\_id & ~3 // for src1 operand  
 posneg ← 0 if negative form, 1 otherwise  
 NFMA\_SS(vl, dest, k1, msrc, regs\_loaded, src\_base, posneg)

#### Intel C/C++ Compiler Intrinsic Equivalent

V4FMADDSS \_\_m128 \_mm\_4fmadd\_ss(\_\_m128, \_\_m128x4, \_\_m128 \*);  
 V4FMADDSS \_\_m128 \_mm\_mask\_4fmadd\_ss(\_\_m128, \_\_mmask8, \_\_m128x4, \_\_m128 \*);  
 V4FMADDSS \_\_m128 \_mm\_maskz\_4fmadd\_ss(\_\_mmask8, \_\_m128, \_\_m128x4, \_\_m128 \*);  
 V4FNMADDSS \_\_m128 \_mm\_4fnmadd\_ss(\_\_m128, \_\_m128x4, \_\_m128 \*);  
 V4FNMADDSS \_\_m128 \_mm\_mask\_4fnmadd\_ss(\_\_m128, \_\_mmask8, \_\_m128x4, \_\_m128 \*);  
 V4FNMADDSS \_\_m128 \_mm\_maskz\_4fnmadd\_ss(\_\_mmask8, \_\_m128, \_\_m128x4, \_\_m128 \*);

#### SIMD Floating-Point Exceptions

Overflow, Underflow, Invalid, Precision, Denormal.

#### Other Exceptions

See Type E2; additionally

#UD                    If the EVEX broadcast bit is set to 1.  
 #UD                    If the MODRM.mod = 0b11.

## VP4DPWSSD – Dot Product of Signed Words with Dword Accumulation (4-iterations)

| Opcode/<br>Instruction  | Op/<br>En | 64/32<br>bit Mode<br>Support | CPUID Feature<br>Flag | Description  |
|---|-----------|------------------------------|-----------------------|--|
| EVEX.DDS.512.F2.0F38.W0 52 /r<br>VP4DPWSSD zmm1{k1}{z}, zmm2+3,<br>m128 | A         | V/V                          | AVX512_4VNNIW         | Multiply signed words from source register block indicated by zmm2 by signed words from m128 and accumulate resulting signed dwords in zmm1. |

### Instruction Operand Encoding

| Op/En | Tuple     | Operand 1        | Operand 2     | Operand 3     | Operand 4 |
|-------|-----------|------------------|---------------|---------------|-----------|
| A     | Tuple1_4X | ModRM:reg (r, w) | EVEX.vvvv (r) | ModRM:r/m (r) | NA        |

### Description

This instruction computes 4 sequential register source-block dot-products of two signed word operands with doubleword accumulation; see Figure 7-1 below. The memory operand is sequentially selected in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based on that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if any bit of the lowest 16-bits of the mask is set to 1 or if a “no masking” encoding is used.

The tuple type Tuple1\_4X implies that four 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.




Figure 7-1. Register Source-Block Dot Product of Two Signed Word Operands with Doubleword Accumulation<sup>1</sup>

#### NOTES:

1. For illustration purposes, one source-block dot product instance is shown out of the four.

**Operation**

src\_reg\_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

VP4DPWSSD dest, src1, src2

(KL,VL) = (16,512)

N ← 4

ORIGDEST ← DEST

src\_base ← src\_reg\_id & ~ (N-1) // for src1 operand

FOR i ← 0 to KL-1:

  IF k1[i] or \*no writemask\*:

    FOR m ← 0 to N-1:

      t ← SRC2.dword[m]

      p1dword ← reg[src\_base+m].word[2\*i] \* t.word[0]

      p2dword ← reg[src\_base+m].word[2\*i+1] \* t.word[1]

      DEST.dword[i] ← DEST.dword[i] + p1dword + p2dword

  ELSE IF \*zeroing\*:

    DEST.dword[i] ← 0

  ELSE

    DEST.dword[i] ← ORIGDEST.dword[i]

DEST[MAX\_VL-1:VL] ← 0

**Intel C/C++ Compiler Intrinsic Equivalent**

VP4DPWSSD \_\_m512i\_mm512\_4dpwssd\_epi32(\_\_m512i, \_\_m512ix4, \_\_m128i \*);

VP4DPWSSD \_\_m512i\_mm512\_mask\_4dpwssd\_epi32(\_\_m512i, \_\_mmask16, \_\_m512ix4, \_\_m128i \*);

VP4DPWSSD \_\_m512i\_mm512\_maskz\_4dpwssd\_epi32(\_\_mmask16, \_\_m512i, \_\_m512ix4, \_\_m128i \*);

**SIMD Floating-Point Exceptions**

None.

**Other Exceptions**

See Type E4; additionally

#UD                    If the EVEX broadcast bit is set to 1.

#UD                    If the MODRM.mod = 0b11.

## 8. Updates to Chapter 1, Volume 3A

Change bars show changes to Chapter 1 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide, Part 1*.

-----  
Change to this chapter: Updates to processors covered by manual; added 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series.

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide, Part 1* (order number 253668), the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2* (order number 253669), the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C: System Programming Guide, Part 3* (order number 326019), and the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3D: System Programming Guide, Part 4* (order number 332831) are part of a set that describes the architecture and programming environment of Intel 64 and IA-32 Architecture processors. The other volumes in this set are:

- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture* (order number 253665).
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D: Instruction Set Reference* (order numbers 253666, 253667, 326018 and 334569).
- *The Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers* (order number 335592).

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, describes the basic architecture and programming environment of Intel 64 and IA-32 processors. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D*, describe the instruction set of the processor and the opcode structure. These volumes apply to application programmers and to programmers who write operating systems or executives. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D*, describe the operating-system support environment of Intel 64 and IA-32 processors. These volumes target operating-system and BIOS designers. In addition, *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B*, and *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C* address the programming environment for classes of software that host operating systems. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*, describes the model-specific registers of Intel 64 and IA-32 processors.

## 1.1 INTEL® 64 AND IA-32 PROCESSORS COVERED IN THIS MANUAL

This manual set includes information pertaining primarily to the most recent Intel 64 and IA-32 processors, which include:

- Pentium® processors
- P6 family processors
- Pentium® 4 processors
- Pentium® M processors
- Intel® Xeon® processors
- Pentium® D processors
- Pentium® processor Extreme Editions
- 64-bit Intel® Xeon® processors
- Intel® Core™ Duo processor
- Intel® Core™ Solo processor
- Dual-Core Intel® Xeon® processor LV
- Intel® Core™2 Duo processor
- Intel® Core™2 Quad processor Q6000 series
- Intel® Xeon® processor 3000, 3200 series
- Intel® Xeon® processor 5000 series
- Intel® Xeon® processor 5100, 5300 series

## ABOUT THIS MANUAL

- Intel® Core™2 Extreme processor X7000 and X6800 series
- Intel® Core™2 Extreme QX6000 series
- Intel® Xeon® processor 7100 series
- Intel® Pentium® Dual-Core processor
- Intel® Xeon® processor 7200, 7300 series
- Intel® Core™2 Extreme QX9000 series
- Intel® Xeon® processor 5200, 5400, 7400 series
- Intel® Core™2 Extreme processor QX9000 and X9000 series
- Intel® Core™2 Quad processor Q9000 series
- Intel® Core™2 Duo processor E8000, T9000 series
- Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are built from 45 nm and 32 nm processes.
- Intel® Core™ i7 processor
- Intel® Core™ i5 processor
- Intel® Xeon® processor E7-8800/4800/2800 product families
- Intel® Core™ i7-3930K processor
- 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series
- Intel® Xeon® processor E3-1200 product family
- Intel® Xeon® processor E5-2400/1400 product family
- Intel® Xeon® processor E5-4600/2600/1600 product family
- 3rd generation Intel® Core™ processors
- Intel® Xeon® processor E3-1200 v2 product family
- Intel® Xeon® processor E5-2400/1400 v2 product families
- Intel® Xeon® processor E5-4600/2600/1600 v2 product families
- Intel® Xeon® processor E7-8800/4800/2800 v2 product families
- 4th generation Intel® Core™ processors
- The Intel® Core™ M processor family
- Intel® Core™ i7-59xx Processor Extreme Edition
- Intel® Core™ i7-49xx Processor Extreme Edition
- Intel® Xeon® processor E3-1200 v3 product family
- Intel® Xeon® processor E5-2600/1600 v3 product families
- 5th generation Intel® Core™ processors
- Intel® Xeon® processor D-1500 product family
- Intel® Xeon® processor E5 v4 family
- Intel® Atom™ processor X7-Z8000 and X5-Z8000 series
- Intel® Atom™ processor Z3400 series
- Intel® Atom™ processor Z3500 series
- 6th generation Intel® Core™ processors
- Intel® Xeon® processor E3-1500m v5 product family
- 7th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series
- Intel® Xeon® Processor Scalable Family
- 8th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series

P6 family processors are IA-32 processors based on the P6 family microarchitecture. This includes the Pentium® Pro, Pentium® II, Pentium® III, and Pentium® III Xeon® processors.

The Pentium® 4, Pentium® D, and Pentium® processor Extreme Editions are based on the Intel NetBurst® microarchitecture. Most early Intel® Xeon® processors are based on the Intel NetBurst® microarchitecture. Intel Xeon processor 5000, 7100 series are based on the Intel NetBurst® microarchitecture.

The Intel® Core™ Duo, Intel® Core™ Solo and dual-core Intel® Xeon® processor LV are based on an improved Pentium® M processor microarchitecture.

The Intel® Xeon® processor 3000, 3200, 5100, 5300, 7200, and 7300 series, Intel® Pentium® dual-core, Intel® Core™2 Duo, Intel® Core™2 Quad, and Intel® Core™2 Extreme processors are based on Intel® Core™ microarchitecture.

The Intel® Xeon® processor 5200, 5400, 7400 series, Intel® Core™2 Quad processor Q9000 series, and Intel® Core™2 Extreme processors QX9000, X9000 series, Intel® Core™2 processor E8000 series are based on Enhanced Intel® Core™ microarchitecture.

The Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are based on the Intel® Atom™ microarchitecture and supports Intel 64 architecture.

P6 family, Pentium® M, Intel® Core™ Solo, Intel® Core™ Duo processors, dual-core Intel® Xeon® processor LV, and early generations of Pentium 4 and Intel Xeon processors support IA-32 architecture. The Intel® Atom™ processor Z5xx series support IA-32 architecture.

The Intel® Xeon® processor 3000, 3200, 5000, 5100, 5200, 5300, 5400, 7100, 7200, 7300, 7400 series, Intel® Core™2 Duo, Intel® Core™2 Extreme, Intel® Core™2 Quad processors, Pentium® D processors, Pentium® Dual-Core processor, newer generations of Pentium 4 and Intel Xeon processor family support Intel® 64 architecture.

The Intel® Core™ i7 processor and Intel® Xeon® processor 3400, 5500, 7500 series are based on 45 nm Intel® microarchitecture code name Nehalem. Intel® microarchitecture code name Westmere is a 32 nm version of Intel® microarchitecture code name Nehalem. Intel® Xeon® processor 5600 series, Intel Xeon processor E7 and various Intel Core i7, i5, i3 processors are based on Intel® microarchitecture code name Westmere. These processors support Intel 64 architecture.

The Intel® Xeon® processor E5 family, Intel® Xeon® processor E3-1200 family, Intel® Xeon® processor E7-8800/4800/2800 product families, Intel® Core™ i7-3930K processor, and 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series are based on the Intel® microarchitecture code name Sandy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E7-8800/4800/2800 v2 product families, Intel® Xeon® processor E3-1200 v2 product family and 3rd generation Intel® Core™ processors are based on the Intel® microarchitecture code name Ivy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E5-4600/2600/1600 v2 product families, Intel® Xeon® processor E5-2400/1400 v2 product families and Intel® Core™ i7-49xx Processor Extreme Edition are based on the Intel® microarchitecture code name Ivy Bridge-E and support Intel 64 architecture.

The Intel® Xeon® processor E3-1200 v3 product family and 4th Generation Intel® Core™ processors are based on the Intel® microarchitecture code name Haswell and support Intel 64 architecture.

The Intel® Xeon® processor E5-2600/1600 v3 product families and the Intel® Core™ i7-59xx Processor Extreme Edition are based on the Intel® microarchitecture code name Haswell-E and support Intel 64 architecture.

The Intel® Atom™ processor Z8000 series is based on the Intel microarchitecture code name Airmont.

The Intel® Atom™ processor Z3400 series and the Intel® Atom™ processor Z3500 series are based on the Intel microarchitecture code name Silvermont.

The Intel® Core™ M processor family, 5th generation Intel® Core™ processors, Intel® Xeon® processor D-1500 product family and the Intel® Xeon® processor E5 v4 family are based on the Intel® microarchitecture code name Broadwell and support Intel 64 architecture.

The Intel® Xeon® Processor Scalable Family, Intel® Xeon® processor E3-1500m v5 product family and 6th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Skylake and support Intel 64 architecture.

The 7th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Kaby Lake and support Intel 64 architecture.



The Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series is based on the Intel® microarchitecture code name Knights Landing and supports Intel 64 architecture.

The 8th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Coffee Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series is based on the Intel® microarchitecture code name Knights Mill and supports Intel 64 architecture.

IA-32 architecture is the instruction set architecture and programming environment for Intel's 32-bit microprocessors. Intel® 64 architecture is the instruction set architecture and programming environment which is the superset of Intel's 32-bit and 64-bit architectures. It is compatible with the IA-32 architecture.

## 1.2 OVERVIEW OF THE SYSTEM PROGRAMMING GUIDE

A description of this manual's content follows<sup>1</sup>:

**Chapter 1 — About This Manual.** Gives an overview of all eight volumes of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*. It also describes the notational conventions in these manuals and lists related Intel manuals and documentation of interest to programmers and hardware designers.

**Chapter 2 — System Architecture Overview.** Describes the modes of operation used by Intel 64 and IA-32 processors and the mechanisms provided by the architectures to support operating systems and executives, including the system-oriented registers and data structures and the system-oriented instructions. The steps necessary for switching between real-address and protected modes are also identified.

**Chapter 3 — Protected-Mode Memory Management.** Describes the data structures, registers, and instructions that support segmentation and paging. The chapter explains how they can be used to implement a "flat" (unsegmented) memory model or a segmented memory model.

**Chapter 4 — Paging.** Describes the paging modes supported by Intel 64 and IA-32 processors.

**Chapter 5 — Protection.** Describes the support for page and segment protection provided in the Intel 64 and IA-32 architectures. This chapter also explains the implementation of privilege rules, stack switching, pointer validation, user and supervisor modes.

**Chapter 6 — Interrupt and Exception Handling.** Describes the basic interrupt mechanisms defined in the Intel 64 and IA-32 architectures, shows how interrupts and exceptions relate to protection, and describes how the architecture handles each exception type. Reference information for each exception is given in this chapter. Includes programming the LINT0 and LINT1 inputs and gives an example of how to program the LINT0 and LINT1 pins for specific interrupt vectors.

**Chapter 7 — Task Management.** Describes mechanisms the Intel 64 and IA-32 architectures provide to support multitasking and inter-task protection.

**Chapter 8 — Multiple-Processor Management.** Describes the instructions and flags that support multiple processors with shared memory, memory ordering, and Intel® Hyper-Threading Technology. Includes MP initialization for P6 family processors and gives an example of how to use the MP protocol to boot P6 family processors in an MP system.

**Chapter 9 — Processor Management and Initialization.** Defines the state of an Intel 64 or IA-32 processor after reset initialization. This chapter also explains how to set up an Intel 64 or IA-32 processor for real-address mode operation and protected-mode operation, and how to switch between modes.

**Chapter 10 — Advanced Programmable Interrupt Controller (APIC).** Describes the programming interface to the local APIC and gives an overview of the interface between the local APIC and the I/O APIC. Includes APIC bus message formats and describes the message formats for messages transmitted on the APIC bus for P6 family and Pentium processors.

**Chapter 11 — Memory Cache Control.** Describes the general concept of caching and the caching mechanisms supported by the Intel 64 or IA-32 architectures. This chapter also describes the memory type range registers (MTRRs) and how they can be used to map memory types of physical memory. Information on using the new cache

1. Model-Specific Registers have been moved out of this volume and into a separate volume: *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*.

control and memory streaming instructions introduced with the Pentium III, Pentium 4, and Intel Xeon processors is also given.

**Chapter 12 — Intel® MMX™ Technology System Programming.** Describes those aspects of the Intel® MMX™ technology that must be handled and considered at the system programming level, including: task switching, exception handling, and compatibility with existing system environments.

**Chapter 13 — System Programming For Instruction Set Extensions And Processor Extended States.** Describes the operating system requirements to support SSE/SSE2/SSE3/SSSE3/SSE4 extensions, including task switching, exception handling, and compatibility with existing system environments. The latter part of this chapter describes the extensible framework of operating system requirements to support processor extended states. Processor extended state may be required by instruction set extensions beyond those of SSE/SSE2/SSE3/SSSE3/SSE4 extensions.

**Chapter 14 — Power and Thermal Management.** Describes facilities of Intel 64 and IA-32 architecture used for power management and thermal monitoring.

**Chapter 15 — Machine-Check Architecture.** Describes the machine-check architecture and machine-check exception mechanism found in the Pentium 4, Intel Xeon, and P6 family processors. Additionally, a signaling mechanism for software to respond to hardware corrected machine check error is covered.

**Chapter 16 — Interpreting Machine-Check Error Codes.** Gives an example of how to interpret the error codes for a machine-check error that occurred on a P6 family processor.

**Chapter 17 — Debug, Branch Profile, TSC, and Resource Monitoring Features.** Describes the debugging registers and other debug mechanism provided in Intel 64 or IA-32 processors. This chapter also describes the time-stamp counter.

**Chapter 18 — Performance Monitoring.** Describes the Intel 64 and IA-32 architectures' facilities for monitoring performance.

**Chapter 19 — Performance-Monitoring Events.** Lists architectural performance events. Non-architectural performance events (i.e. model-specific events) are listed for each generation of microarchitecture.

**Chapter 20 — 8086 Emulation.** Describes the real-address and virtual-8086 modes of the IA-32 architecture.

**Chapter 21 — Mixing 16-Bit and 32-Bit Code.** Describes how to mix 16-bit and 32-bit code modules within the same program or task.

**Chapter 22 — IA-32 Architecture Compatibility.** Describes architectural compatibility among IA-32 processors.

**Chapter 23 — Introduction to Virtual Machine Extensions.** Describes the basic elements of virtual machine architecture and the virtual machine extensions for Intel 64 and IA-32 Architectures.

**Chapter 24 — Virtual Machine Control Structures.** Describes components that manage VMX operation. These include the working-VMCS pointer and the controlling-VMCS pointer.

**Chapter 25 — VMX Non-Root Operation.** Describes the operation of a VMX non-root operation. Processor operation in VMX non-root mode can be restricted programmatically such that certain operations, events or conditions can cause the processor to transfer control from the guest (running in VMX non-root mode) to the monitor software (running in VMX root mode).

**Chapter 26 — VM Entries.** Describes VM entries. VM entry transitions the processor from the VMM running in VMX root-mode to a VM running in VMX non-root mode. VM-Entry is performed by the execution of VMLAUNCH or VMRESUME instructions.

**Chapter 27 — VM Exits.** Describes VM exits. Certain events, operations or situations while the processor is in VMX non-root operation may cause VM-exit transitions. In addition, VM exits can also occur on failed VM entries.

**Chapter 28 — VMX Support for Address Translation.** Describes virtual-machine extensions that support address translation and the virtualization of physical memory.

**Chapter 29 — APIC Virtualization and Virtual Interrupts.** Describes the VMCS including controls that enable the virtualization of interrupts and the Advanced Programmable Interrupt Controller (APIC).

**Chapter 30 — VMX Instruction Reference.** Describes the virtual-machine extensions (VMX). VMX is intended for a system executive to support virtualization of processor hardware and a system software layer acting as a host to multiple guest software environments.

**Chapter 31 — Virtual-Machine Monitor Programming Considerations.** Describes programming considerations for VMMs. VMMs manage virtual machines (VMs).

**Chapter 32 — Virtualization of System Resources.** Describes the virtualization of the system resources. These include: debugging facilities, address translation, physical memory, and microcode update facilities.

**Chapter 33 — Handling Boundary Conditions in a Virtual Machine Monitor.** Describes what a VMM must consider when handling exceptions, interrupts, error conditions, and transitions between activity states.

**Chapter 34 — System Management Mode.** Describes Intel 64 and IA-32 architectures' system management mode (SMM) facilities.

**Chapter 35 — Intel® Processor Trace.** Describes details of Intel® Processor Trace.

**Chapter 36 — Introduction to Intel® Software Guard Extensions.** Provides an overview of the Intel® Software Guard Extensions (Intel® SGX) set of instructions.

**Chapter 37 — Enclave Access Control and Data Structures.** Describes Enclave Access Control procedures and defines various Intel SGX data structures.

**Chapter 38 — Enclave Operation.** Describes enclave creation and initialization, adding pages and measuring an enclave, and enclave entry and exit.

**Chapter 39 — Enclave Exiting Events.** Describes enclave-exiting events (EEE) and asynchronous enclave exit (AEX).

**Chapter 40 — SGX Instruction References.** Describes the supervisor and user level instructions provided by Intel SGX.

**Chapter 41 — Intel® SGX Interactions with IA32 and Intel® 64 Architecture.** Describes the Intel SGX collection of enclave instructions for creating protected execution environments on processors supporting IA32 and Intel 64 architectures.

**Chapter 42 — Enclave Code Debug and Profiling.** Describes enclave code debug processes and options.

**Appendix A — VMX Capability Reporting Facility.** Describes the VMX capability MSRs. Support for specific VMX features is determined by reading capability MSRs.

**Appendix B — Field Encoding in VMCS.** Enumerates all fields in the VMCS and their encodings. Fields are grouped by width (16-bit, 32-bit, etc.) and type (guest-state, host-state, etc.).

**Appendix C — VM Basic Exit Reasons.** Describes the 32-bit fields that encode reasons for a VM exit. Examples of exit reasons include, but are not limited to: software interrupts, processor exceptions, software traps, NMIs, external interrupts, and triple faults.

## 1.3 NOTATIONAL CONVENTIONS

This manual uses specific notation for data-structure formats, for symbolic representation of instructions, and for hexadecimal and binary numbers. A review of this notation makes the manual easier to read.

### 1.3.1 Bit and Byte Order

In illustrations of data structures in memory, smaller addresses appear toward the bottom of the figure; addresses increase toward the top. Bit positions are numbered from right to left. The numerical value of a set bit is equal to two raised to the power of the bit position. Intel 64 and IA-32 processors are "little endian" machines; this means the bytes of a word are numbered starting from the least significant byte. Figure 1-1 illustrates these conventions.

### 1.3.2 Reserved Bits and Software Compatibility

In many register and memory layout descriptions, certain bits are marked as reserved. When bits are marked as reserved, it is essential for compatibility with future processors that software treat these bits as having a future, though unknown, effect. The behavior of reserved bits should be regarded as not only undefined, but unpredictable. Software should follow these guidelines in dealing with reserved bits:

- Do not depend on the states of any reserved bits when testing the values of registers which contain such bits. Mask out the reserved bits before testing.
- Do not depend on the states of any reserved bits when storing to memory or to a register.
- Do not depend on the ability to retain information written into any reserved bits.
- When loading a register, always load the reserved bits with the values indicated in the documentation, if any, or reload them with values previously read from the same register.

**NOTE**

Avoid any software dependence upon the state of reserved bits in Intel 64 and IA-32 registers. Depending upon the values of reserved register bits will make software dependent upon the unspecified manner in which the processor handles these bits. Programs that depend upon reserved values risk incompatibility with future processors.

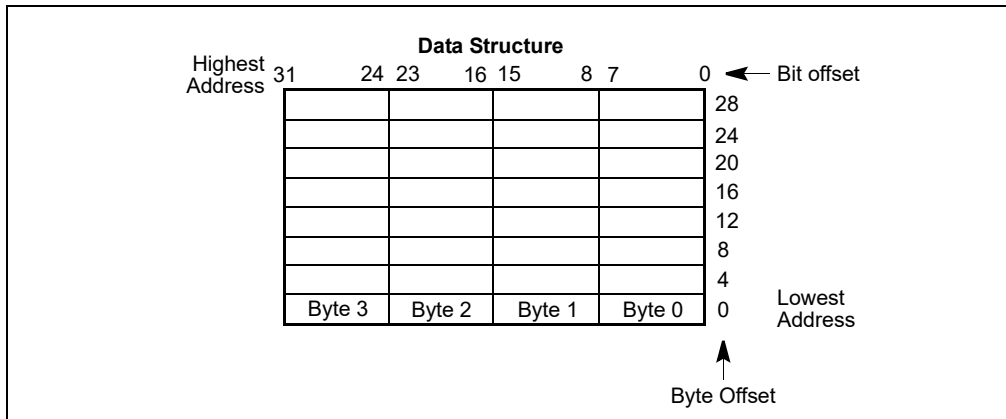


Figure 1-1. Bit and Byte Order

### 1.3.3 Instruction Operands

When instructions are represented symbolically, a subset of assembly language is used. In this subset, an instruction has the following format:

```
label: mnemonic argument1, argument2, argument3
```

where:

- A **label** is an identifier which is followed by a colon.
- A **mnemonic** is a reserved name for a class of instruction opcodes which have the same function.
- The operands **argument1**, **argument2**, and **argument3** are optional. There may be from zero to three operands, depending on the opcode. When present, they take the form of either literals or identifiers for data items. Operand identifiers are either reserved names of registers or are assumed to be assigned to data items declared in another part of the program (which may not be shown in the example).

When two operands are present in an arithmetic or logical instruction, the right operand is the source and the left operand is the destination.

For example:

```
LOADREG: MOV EAX, SUBTOTAL
```

In this example **LOADREG** is a label, **MOV** is the mnemonic identifier of an opcode, **EAX** is the destination operand, and **SUBTOTAL** is the source operand. Some assembly languages put the source and destination in reverse order.

### 1.3.4 Hexadecimal and Binary Numbers

Base 16 (hexadecimal) numbers are represented by a string of hexadecimal digits followed by the character **H** (for example, **F82EH**). A hexadecimal digit is a character from the following set: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, and F.

Base 2 (binary) numbers are represented by a string of 1s and 0s, sometimes followed by the character **B** (for example, **1010B**). The "B" designation is only used in situations where confusion as to the type of number might arise.

### 1.3.5 Segmented Addressing

The processor uses byte addressing. This means memory is organized and accessed as a sequence of bytes. Whether one or more bytes are being accessed, a byte address is used to locate the byte or bytes memory. The range of memory that can be addressed is called an **address space**.

The processor also supports segmented addressing. This is a form of addressing where a program may have many independent address spaces, called **segments**. For example, a program can keep its code (instructions) and stack in separate segments. Code addresses would always refer to the code space, and stack addresses would always refer to the stack space. The following notation is used to specify a byte address within a segment:

```
Segment-register:Byte-address
```

For example, the following segment address identifies the byte at address **FF79H** in the segment pointed by the **DS** register:

```
DS:FF79H
```

The following segment address identifies an instruction address in the code segment. The **CS** register points to the code segment and the **EIP** register contains the address of the instruction.

```
CS:EIP
```

### 1.3.6 Syntax for CPUID, CR, and MSR Values

Obtain feature flags, status, and system information by using the CPUID instruction, by checking control register bits, and by reading model-specific registers. We are moving toward a single syntax to represent this type of information. See Figure 1-2.




Figure 1-2. Syntax for CPUID, CR, and MSR Data Presentation

### 1.3.7 Exceptions

An exception is an event that typically occurs when an instruction causes an error. For example, an attempt to divide by zero generates an exception. However, some exceptions, such as breakpoints, occur under other conditions. Some types of exceptions may provide error codes. An error code reports additional information about the error. An example of the notation used to show an exception and error code is shown below:

#PF(fault code)

This example refers to a page-fault exception under conditions where an error code naming a type of fault is reported. Under some conditions, exceptions which produce error codes may not be able to report an accurate code. In this case, the error code is zero, as shown below for a general-protection exception:

#GP(0)

## 1.4 RELATED LITERATURE

Literature related to Intel 64 and IA-32 processors is listed and viewable on-line at:

<http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>

See also:

- The data sheet for a particular Intel 64 or IA-32 processor
- The specification update for a particular Intel 64 or IA-32 processor
- Intel® C++ Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Fortran Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Software Development Tools:  
<https://software.intel.com/en-us/intel-sdp-home>
- Intel® 64 and IA-32 Architectures Software Developer's Manual (in one, four or ten volumes):  
<https://software.intel.com/en-us/articles/intel-sdm>
- Intel® 64 and IA-32 Architectures Optimization Reference Manual:  
<https://software.intel.com/en-us/articles/intel-sdm#optimization>
- Intel 64 Architecture x2APIC Specification:  
<http://www.intel.com/content/www/us/en/architecture-and-technology/64-architecture-x2apic-specification.html>
- Intel® Trusted Execution Technology Measured Launched Environment Programming Guide:  
<http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>
- Developing Multi-threaded Applications: A Platform Consistent Approach:  
<https://software.intel.com/sites/default/files/article/147714/51534-developing-multithreaded-applications.pdf>
- Using Spin-Loops on Intel® Pentium® 4 Processor and Intel® Xeon® Processor:  
<https://software.intel.com/sites/default/files/22/30/25602>
- Performance Monitoring Unit Sharing Guide  
<http://software.intel.com/file/30388>

Literature related to selected features in future Intel processors are available at:

- Intel® Architecture Instruction Set Extensions Programming Reference  
<https://software.intel.com/en-us/isa-extensions>
- Intel® Software Guard Extensions (Intel® SGX) Programming Reference  
<https://software.intel.com/en-us/isa-extensions/intel-sgx>

More relevant links are:

- Intel® Developer Zone:  
<https://software.intel.com/en-us>
- Developer centers:  
<http://www.intel.com/content/www/us/en/hardware-developers/developer-centers.html>
- Processor support general link:  
<http://www.intel.com/support/processors/>
- Intel® Hyper-Threading Technology (Intel® HT Technology):  
<http://www.intel.com/technology/platform-technology/hyper-threading/index.htm>

## 9. Updates to Chapter 4, Volume 3A

Change bars show changes to Chapter 4 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide, Part 1*.

-----  
Changes to this chapter: Correction to text alignment in Figure 4-4 "Formats of CR3 and Paging-Structure Entries with 32-Bit Paging".



Chapter 3 explains how segmentation converts logical addresses to linear addresses. **Paging** (or linear-address translation) is the process of translating linear addresses so that they can be used to access memory or I/O devices. Paging translates each linear address to a **physical address** and determines, for each translation, what accesses to the linear address are allowed (the address's **access rights**) and the type of caching used for such accesses (the address's **memory type**).

Intel-64 processors support three different paging modes. These modes are identified and defined in Section 4.1. Section 4.2 gives an overview of the translation mechanism that is used in all modes. Section 4.3, Section 4.4, and Section 4.5 discuss the three paging modes in detail.

Section 4.6 details how paging determines and uses access rights. Section 4.7 discusses exceptions that may be generated by paging (page-fault exceptions). Section 4.8 considers data which the processor writes in response to linear-address accesses (accessed and dirty flags).

Section 4.9 describes how paging determines the memory types used for accesses to linear addresses. Section 4.10 provides details of how a processor may cache information about linear-address translation. Section 4.11 outlines interactions between paging and certain VMX features. Section 4.12 gives an overview of how paging can be used to implement virtual memory.

## 4.1 PAGING MODES AND CONTROL BITS

Paging behavior is controlled by the following control bits:

- The WP and PG flags in control register CR0 (bit 16 and bit 31, respectively).
- The PSE, PAE, PGE, PCIDE, SMEP, SMAP, and PKE flags in control register CR4 (bit 4, bit 5, bit 7, bit 17, bit 20, bit 21, and bit 22, respectively).
- The LME and NXE flags in the IA32\_EFER MSR (bit 8 and bit 11, respectively).
- The AC flag in the EFLAGS register (bit 18).

Software enables paging by using the MOV to CR0 instruction to set CR0.PG. Before doing so, software should ensure that control register CR3 contains the physical address of the first paging structure that the processor will use for linear-address translation (see Section 4.2) and that structure is initialized as desired. See Table 4-3, Table 4-7, and Table 4-12 for the use of CR3 in the different paging modes.

Section 4.1.1 describes how the values of CR0.PG, CR4.PAE, and IA32\_EFER.LME determine whether paging is in use and, if so, which of three paging modes is in use. Section 4.1.2 explains how to manage these bits to establish or make changes in paging modes. Section 4.1.3 discusses how CR0.WP, CR4.PSE, CR4.PGE, CR4.PCIDE, CR4.SMEP, CR4.SMAP, CR4.PKE, and IA32\_EFER.NXE modify the operation of the different paging modes.

### 4.1.1 Three Paging Modes

If CR0.PG = 0, paging is not used. The logical processor treats all linear addresses as if they were physical addresses. CR4.PAE and IA32\_EFER.LME are ignored by the processor, as are CR0.WP, CR4.PSE, CR4.PGE, CR4.SMEP, CR4.SMAP, and IA32\_EFER.NXE.

Paging is enabled if CR0.PG = 1. Paging can be enabled only if protection is enabled (CR0.PE = 1). If paging is enabled, one of three paging modes is used. The values of CR4.PAE and IA32\_EFER.LME determine which paging mode is used:

- If CR0.PG = 1 and CR4.PAE = 0, **32-bit paging** is used. 32-bit paging is detailed in Section 4.3. 32-bit paging uses CR0.WP, CR4.PSE, CR4.PGE, CR4.SMEP, and CR4.SMAP as described in Section 4.1.3.
- If CR0.PG = 1, CR4.PAE = 1, and IA32\_EFER.LME = 0, **PAE paging** is used. PAE paging is detailed in Section 4.4. PAE paging uses CR0.WP, CR4.PGE, CR4.SMEP, CR4.SMAP, and IA32\_EFER.NXE as described in Section 4.1.3.

## PAGING

- If CR0.PG = 1, CR4.PAE = 1, and IA32\_EFER.LME = 1, 4-level paging<sup>1</sup> is used.<sup>2</sup> 4-level paging is detailed in Section 4.5. 4-level paging uses CR0.WP, CR4.PGE, CR4.PCIDE, CR4.SMEP, CR4.SMAP, CR4.PKE, and IA32\_EFER.NXE as described in Section 4.1.3. 4-level paging is available only on processors that support the Intel 64 architecture.

The three paging modes differ with regard to the following details:

- Linear-address width. The size of the linear addresses that can be translated.
- Physical-address width. The size of the physical addresses produced by paging.
- Page size. The granularity at which linear addresses are translated. Linear addresses on the same page are translated to corresponding physical addresses on the same page.
- Support for execute-disable access rights. In some paging modes, software can be prevented from fetching instructions from pages that are otherwise readable.
- Support for PCIDs. With 4-level paging, software can enable a facility by which a logical processor caches information for multiple linear-address spaces. The processor may retain cached information when software switches between different linear-address spaces.
- Support for protection keys. With 4-level paging, software can enable a facility by which each linear address is associated with a protection key. Software can use a new control register to determine, for each protection key, how software can access linear addresses associated with that protection key.

Table 4-1 illustrates the principal differences between the three paging modes.

**Table 4-1. Properties of Different Paging Modes**

| Paging Mode | PG in CR0 | PAE in CR4 | LME in IA32_EFER | Lin.-Addr. Width | Phys.-Addr. Width <sup>1</sup> | Page Sizes                        | Supports Execute-Disable? | Supports PCIDs and protection keys? |
|-------------|-----------|------------|------------------|------------------|--------------------------------|-----------------------------------|---------------------------|-------------------------------------|
| None        | 0         | N/A        | N/A              | 32               | 32                             | N/A                               | No                        | No                                  |
| 32-bit      | 1         | 0          | 0 <sup>2</sup>   | 32               | Up to 40 <sup>3</sup>          | 4 KB<br>4 MB <sup>4</sup>         | No                        | No                                  |
| PAE         | 1         | 1          | 0                | 32               | Up to 52                       | 4 KB<br>2 MB                      | Yes <sup>5</sup>          | No                                  |
| 4-level     | 1         | 1          | 1                | 48               | Up to 52                       | 4 KB<br>2 MB<br>1 GB <sup>6</sup> | Yes <sup>5</sup>          | Yes <sup>7</sup>                    |

### NOTES:

1. The physical-address width is always bounded by MAXPHYADDR; see Section 4.1.4.
2. The processor ensures that IA32\_EFER.LME must be 0 if CR0.PG = 1 and CR4.PAE = 0.
3. 32-bit paging supports physical-address widths of more than 32 bits only for 4-MByte pages and only if the PSE-36 mechanism is supported; see Section 4.1.4 and Section 4.3.
4. 4-MByte pages are used with 32-bit paging only if CR4.PSE = 1; see Section 4.3.
5. Execute-disable access rights are applied only if IA32\_EFER.NXE = 1; see Section 4.6.
6. Not all processors that support 4-level paging support 1-GByte pages; see Section 4.1.4.
7. PCIDs are used only if CR4.PCIDE = 1; see Section 4.10.1. Protection keys are used only if certain conditions hold; see Section 4.6.2.

- 
1. Earlier versions of this manual used the term "IA-32e paging" to identify 4-level paging.
  2. The LMA flag in the IA32\_EFER MSR (bit 10) is a status bit that indicates whether the logical processor is in IA-32e mode (and thus using 4-level paging). The processor always sets IA32\_EFER.LMA to CR0.PG & IA32\_EFER.LME. Software cannot directly modify IA32\_EFER.LMA; an execution of WRMSR to the IA32\_EFER MSR ignores bit 10 of its source operand.

Because they are used only if IA32\_EFER.LME = 0, 32-bit paging and PAE paging are used only in legacy protected mode. Because legacy protected mode cannot produce linear addresses larger than 32 bits, 32-bit paging and PAE paging translate 32-bit linear addresses.

Because it is used only if IA32\_EFER.LME = 1, 4-level paging is used only in IA-32e mode. (In fact, it is the use of 4-level paging that defines IA-32e mode.) IA-32e mode has two sub-modes:

- Compatibility mode. This mode uses only 32-bit linear addresses. 4-level paging treats bits 47:32 of such an address as all 0.
- 64-bit mode. While this mode produces 64-bit linear addresses, the processor ensures that bits 63:47 of such an address are identical.<sup>1</sup> 4-level paging does not use bits 63:48 of such addresses.

### 4.1.2 Paging-Mode Enabling

If CR0.PG = 1, a logical processor is in one of three paging modes, depending on the values of CR4.PAE and IA32\_EFER.LME. Figure 4-1 illustrates how software can enable these modes and make transitions between them. The following items identify certain limitations and other details:

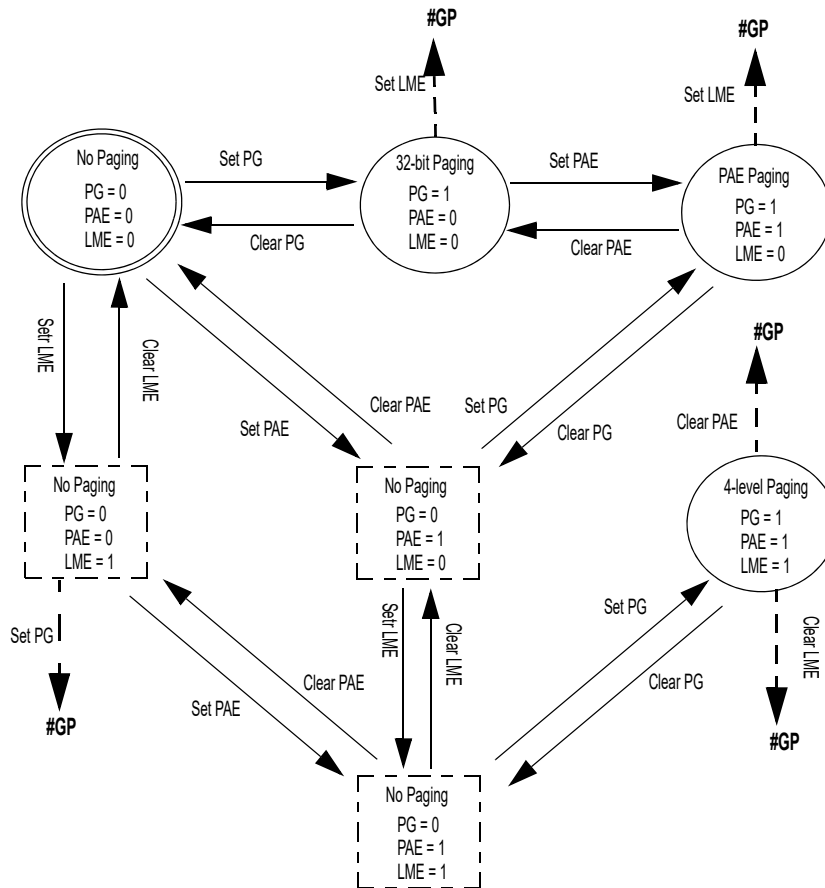


Figure 4-1. Enabling and Changing Paging Modes

- IA32\_EFER.LME cannot be modified while paging is enabled (CR0.PG = 1). Attempts to do so using WRMSR cause a general-protection exception (#GP(0)).

1. Such an address is called **canonical**. Use of a non-canonical linear address in 64-bit mode produces a general-protection exception (#GP(0)); the processor does not attempt to translate non-canonical linear addresses using 4-level paging.

- Paging cannot be enabled (by setting CR0.PG to 1) while CR4.PAE = 0 and IA32\_EFER.LME = 1. Attempts to do so using MOV to CR0 cause a general-protection exception (#GP(0)).
- CR4.PAE cannot be cleared while 4-level paging is active (CR0.PG = 1 and IA32\_EFER.LME = 1). Attempts to do so using MOV to CR4 cause a general-protection exception (#GP(0)).
- Regardless of the current paging mode, software can disable paging by clearing CR0.PG with MOV to CR0.<sup>1</sup>
- Software can make transitions between 32-bit paging and PAE paging by changing the value of CR4.PAE with MOV to CR4.
- Software cannot make transitions directly between 4-level paging and either of the other two paging modes. It must first disable paging (by clearing CR0.PG with MOV to CR0), then set CR4.PAE and IA32\_EFER.LME to the desired values (with MOV to CR4 and WRMSR), and then re-enable paging (by setting CR0.PG with MOV to CR0). As noted earlier, an attempt to clear either CR4.PAE or IA32\_EFER.LME cause a general-protection exception (#GP(0)).
- VMX transitions allow transitions between paging modes that are not possible using MOV to CR or WRMSR. This is because VMX transitions can load CR0, CR4, and IA32\_EFER in one operation. See Section 4.11.1.

### 4.1.3 Paging-Mode Modifiers

Details of how each paging mode operates are determined by the following control bits:

- The WP flag in CR0 (bit 16).
- The PSE, PGE, PCIDE, SMEP, SMAP, and PKE flags in CR4 (bit 4, bit 7, bit 17, bit 20, bit 21, and bit 22 respectively).
- The NXE flag in the IA32\_EFER MSR (bit 11).

CR0.WP allows pages to be protected from supervisor-mode writes. If CR0.WP = 0, supervisor-mode write accesses are allowed to linear addresses with read-only access rights; if CR0.WP = 1, they are not. (User-mode write accesses are never allowed to linear addresses with read-only access rights, regardless of the value of CR0.WP.) Section 4.6 explains how access rights are determined, including the definition of supervisor-mode and user-mode accesses.

CR4.PSE enables 4-MByte pages for 32-bit paging. If CR4.PSE = 0, 32-bit paging can use only 4-KByte pages; if CR4.PSE = 1, 32-bit paging can use both 4-KByte pages and 4-MByte pages. See Section 4.3 for more information. (PAE paging and 4-level paging can use multiple page sizes regardless of the value of CR4.PSE.)

CR4.PGE enables global pages. If CR4.PGE = 0, no translations are shared across address spaces; if CR4.PGE = 1, specified translations may be shared across address spaces. See Section 4.10.2.4 for more information.

CR4.PCIDE enables process-context identifiers (PCIDs) for 4-level paging (CR4.PCIDE can be 1 only when 4-level paging is in use). PCIDs allow a logical processor to cache information for multiple linear-address spaces. See Section 4.10.1 for more information.

CR4.SMEP allows pages to be protected from supervisor-mode instruction fetches. If CR4.SMEP = 1, software operating in supervisor mode cannot fetch instructions from linear addresses that are accessible in user mode. Section 4.6 explains how access rights are determined, including the definition of supervisor-mode accesses and user-mode accessibility.

CR4.SMAP allows pages to be protected from supervisor-mode data accesses. If CR4.SMAP = 1, software operating in supervisor mode cannot access data at linear addresses that are accessible in user mode. Software can override this protection by setting EFLAGS.AC. Section 4.6 explains how access rights are determined, including the definition of supervisor-mode accesses and user-mode accessibility.

CR4.PKE allows each linear address to be associated with a **protection key**. The PKRU register specifies, for each protection key, whether linear addresses with that protection key can be read or written by software. See Section 4.6 for more information.

IA32\_EFER.NXE enables execute-disable access rights for PAE paging and 4-level paging. If IA32\_EFER.NXE = 1, instruction fetches can be prevented from specified linear addresses (even if data reads from the addresses are

---

1. If CR4.PCIDE = 1, an attempt to clear CR0.PG causes a general-protection exception (#GP); software should clear CR4.PCIDE before attempting to disable paging.

allowed). Section 4.6 explains how access rights are determined. (IA32\_EFER.NXE has no effect with 32-bit paging. Software that wants to use this feature to limit instruction fetches from readable pages must use either PAE paging or 4-level paging.)

#### 4.1.4 Enumeration of Paging Features by CPUID

Software can discover support for different paging features using the CPUID instruction:

- PSE: page-size extensions for 32-bit paging.  
If CPUID.01H:EDX.PSE [bit 3] = 1, CR4.PSE may be set to 1, enabling support for 4-MByte pages with 32-bit paging (see Section 4.3).
- PAE: physical-address extension.  
If CPUID.01H:EDX.PAE [bit 6] = 1, CR4.PAE may be set to 1, enabling PAE paging (this setting is also required for 4-level paging).
- PGE: global-page support.  
If CPUID.01H:EDX.PGE [bit 13] = 1, CR4.PGE may be set to 1, enabling the global-page feature (see Section 4.10.2.4).
- PAT: page-attribute table.  
If CPUID.01H:EDX.PAT [bit 16] = 1, the 8-entry page-attribute table (PAT) is supported. When the PAT is supported, three bits in certain paging-structure entries select a memory type (used to determine type of caching used) from the PAT (see Section 4.9.2).
- PSE-36: page-size extensions with 40-bit physical-address extension.  
If CPUID.01H:EDX.PSE-36 [bit 17] = 1, the PSE-36 mechanism is supported, indicating that translations using 4-MByte pages with 32-bit paging may produce physical addresses with up to 40 bits (see Section 4.3).
- PCID: process-context identifiers.  
If CPUID.01H:ECX.PCID [bit 17] = 1, CR4.PCIDE may be set to 1, enabling process-context identifiers (see Section 4.10.1).
- SMEP: supervisor-mode execution prevention.  
If CPUID.(EAX=07H,ECX=0H):EBX.SMEP [bit 7] = 1, CR4.SMEP may be set to 1, enabling supervisor-mode execution prevention (see Section 4.6).
- SMAP: supervisor-mode access prevention.  
If CPUID.(EAX=07H,ECX=0H):EBX.SMAP [bit 20] = 1, CR4.SMAP may be set to 1, enabling supervisor-mode access prevention (see Section 4.6).
- PKU: protection keys.  
If CPUID.(EAX=07H,ECX=0H):ECX.PKU [bit 3] = 1, CR4.PKE may be set to 1, enabling protection keys (see Section 4.6).
- NX: execute disable.  
If CPUID.80000001H:EDX.NX [bit 20] = 1, IA32\_EFER.NXE may be set to 1, allowing PAE paging and 4-level paging to disable execute access to selected pages (see Section 4.6). (Processors that do not support CPUID function 80000001H do not allow IA32\_EFER.NXE to be set to 1.)
- Page1GB: 1-GByte pages.  
If CPUID.80000001H:EDX.Page1GB [bit 26] = 1, 1-GByte pages are supported with 4-level paging (see Section 4.5).
- LM: IA-32e mode support.  
If CPUID.80000001H:EDX.LM [bit 29] = 1, IA32\_EFER.LME may be set to 1, enabling 4-level paging. (Processors that do not support CPUID function 80000001H do not allow IA32\_EFER.LME to be set to 1.)
- CPUID.80000008H:EAX[7:0] reports the physical-address width supported by the processor. (For processors that do not support CPUID function 80000008H, the width is generally 36 if CPUID.01H:EDX.PAE [bit 6] = 1 and 32 otherwise.) This width is referred to as MAXPHYADDR. MAXPHYADDR is at most 52.
- CPUID.80000008H:EAX[15:8] reports the linear-address width supported by the processor. Generally, this value is 48 if CPUID.80000001H:EDX.LM [bit 29] = 1 and 32 otherwise. (Processors that do not support CPUID function 80000008H, support a linear-address width of 32.)

## 4.2 HIERARCHICAL PAGING STRUCTURES: AN OVERVIEW

All three paging modes translate linear addresses using **hierarchical paging structures**. This section provides an overview of their operation. Section 4.3, Section 4.4, and Section 4.5 provide details for the three paging modes.

Every paging structure is 4096 Bytes in size and comprises a number of individual entries. With 32-bit paging, each entry is 32 bits (4 bytes); there are thus 1024 entries in each structure. With PAE paging and 4-level paging, each entry is 64 bits (8 bytes); there are thus 512 entries in each structure. (PAE paging includes one exception, a paging structure that is 32 bytes in size, containing 4 64-bit entries.)

The processor uses the upper portion of a linear address to identify a series of paging-structure entries. The last of these entries identifies the physical address of the region to which the linear address translates (called the **page frame**). The lower portion of the linear address (called the **page offset**) identifies the specific address within that region to which the linear address translates.

Each paging-structure entry contains a physical address, which is either the address of another paging structure or the address of a page frame. In the first case, the entry is said to reference the other paging structure; in the latter, the entry is said to **map a page**.

The first paging structure used for any translation is located at the physical address in CR3. A linear address is translated using the following iterative procedure. A portion of the linear address (initially the uppermost bits) selects an entry in a paging structure (initially the one located using CR3). If that entry references another paging structure, the process continues with that paging structure and with the portion of the linear address immediately below that just used. If instead the entry maps a page, the process completes: the physical address in the entry is that of the page frame and the remaining lower portion of the linear address is the page offset.

The following items give an example for each of the three paging modes (each example locates a 4-KByte page frame):

- With 32-bit paging, each paging structure comprises  $1024 = 2^{10}$  entries. For this reason, the translation process uses 10 bits at a time from a 32-bit linear address. Bits 31:22 identify the first paging-structure entry and bits 21:12 identify a second. The latter identifies the page frame. Bits 11:0 of the linear address are the page offset within the 4-KByte page frame. (See Figure 4-2 for an illustration.)
- With PAE paging, the first paging structure comprises only  $4 = 2^2$  entries. Translation thus begins by using bits 31:30 from a 32-bit linear address to identify the first paging-structure entry. Other paging structures comprise  $512 = 2^9$  entries, so the process continues by using 9 bits at a time. Bits 29:21 identify a second paging-structure entry and bits 20:12 identify a third. This last identifies the page frame. (See Figure 4-5 for an illustration.)
- With 4-level paging, each paging structure comprises  $512 = 2^9$  entries and translation uses 9 bits at a time from a 48-bit linear address. Bits 47:39 identify the first paging-structure entry, bits 38:30 identify a second, bits 29:21 a third, and bits 20:12 identify a fourth. Again, the last identifies the page frame. (See Figure 4-8 for an illustration.)

The translation process in each of the examples above completes by identifying a page frame; the page frame is part of the **translation** of the original linear address. In some cases, however, the paging structures may be configured so that the translation process terminates before identifying a page frame. This occurs if the process encounters a paging-structure entry that is marked “not present” (because its P flag — bit 0 — is clear) or in which a reserved bit is set. In this case, there is no translation for the linear address; an access to that address causes a page-fault exception (see Section 4.7).

In the examples above, a paging-structure entry maps a page with a 4-KByte page frame when only 12 bits remain in the linear address; entries identified earlier always reference other paging structures. That may not apply in other cases. The following items identify when an entry maps a page and when it references another paging structure:

- If more than 12 bits remain in the linear address, bit 7 (PS — page size) of the current paging-structure entry is consulted. If the bit is 0, the entry references another paging structure; if the bit is 1, the entry maps a page.
- If only 12 bits remain in the linear address, the current paging-structure entry always maps a page (bit 7 is used for other purposes).

If a paging-structure entry maps a page when more than 12 bits remain in the linear address, the entry identifies a page frame larger than 4 KBytes. For example, 32-bit paging uses the upper 10 bits of a linear address to locate the first paging-structure entry; 22 bits remain. If that entry maps a page, the page frame is  $2^{22}$  Bytes = 4 MBytes.

32-bit paging supports 4-MByte pages if CR4.PSE = 1. PAE paging and 4-level paging support 2-MByte pages (regardless of the value of CR4.PSE). 4-level paging may support 1-GByte pages (see Section 4.1.4).

Paging structures are given different names based on their uses in the translation process. Table 4-2 gives the names of the different paging structures. It also provides, for each structure, the source of the physical address used to locate it (CR3 or a different paging-structure entry); the bits in the linear address used to select an entry from the structure; and details of whether and how such an entry can map a page.

**Table 4-2. Paging Structures in the Different Paging Modes**

| Paging Structure             | Entry Name | Paging Mode  | Physical Address of Structure | Bits Selecting Entry | Page Mapping                      |
|------------------------------|------------|--------------|-------------------------------|----------------------|-----------------------------------|
| PML4 table                   | PML4E      | 32-bit, PAE  | N/A                           |                      |                                   |
|                              |            | 4-level      | CR3                           | 47:39                | N/A (PS must be 0)                |
| Page-directory-pointer table | PDPTE      | 32-bit       | N/A                           |                      |                                   |
|                              |            | PAE          | CR3                           | 31:30                | N/A (PS must be 0)                |
|                              |            | 4-level      | PML4E                         | 38:30                | 1-GByte page if PS=1 <sup>1</sup> |
| Page directory               | PDE        | 32-bit       | CR3                           | 31:22                | 4-MByte page if PS=1 <sup>2</sup> |
|                              |            | PAE, 4-level | PDPTE                         | 29:21                | 2-MByte page if PS=1              |
| Page table                   | PTE        | 32-bit       | PDE                           | 21:12                | 4-KByte page                      |
|                              |            | PAE, 4-level |                               | 20:12                | 4-KByte page                      |

#### NOTES:

1. Not all processors allow the PS flag to be 1 in PDPTEs; see Section 4.1.4 for how to determine whether 1-GByte pages are supported.
2. 32-bit paging ignores the PS flag in a PDE (and uses the entry to reference a page table) unless CR4.PSE = 1. Not all processors allow CR4.PSE to be 1; see Section 4.1.4 for how to determine whether 4-MByte pages are supported with 32-bit paging.

## 4.3 32-BIT PAGING

A logical processor uses 32-bit paging if CR0.PG = 1 and CR4.PAE = 0. 32-bit paging translates 32-bit linear addresses to 40-bit physical addresses.<sup>1</sup> Although 40 bits corresponds to 1 TByte, linear addresses are limited to 32 bits; at most 4 GBytes of linear-address space may be accessed at any given time.

32-bit paging uses a hierarchy of paging structures to produce a translation for a linear address. CR3 is used to locate the first paging-structure, the page directory. Table 4-3 illustrates how CR3 is used with 32-bit paging.

32-bit paging may map linear addresses to either 4-KByte pages or 4-MByte pages. Figure 4-2 illustrates the translation process when it uses a 4-KByte page; Figure 4-3 covers the case of a 4-MByte page. The following items describe the 32-bit paging process in more detail as well as how the page size is determined:

- A 4-KByte naturally aligned page directory is located at the physical address specified in bits 31:12 of CR3 (see Table 4-3). A page directory comprises 1024 32-bit entries (PDEs). A PDE is selected using the physical address defined as follows:
  - Bits 39:32 are all 0.
  - Bits 31:12 are from CR3.

1. Bits in the range 39:32 are 0 in any physical address used by 32-bit paging except those used to map 4-MByte pages. If the processor does not support the PSE-36 mechanism, this is true also for physical addresses used to map 4-MByte pages. If the processor does support the PSE-36 mechanism and MAXPHYADDR < 40, bits in the range 39:MAXPHYADDR are 0 in any physical address used to map a 4-MByte page. (The corresponding bits are reserved in PDEs.) See Section 4.1.4 for how to determine MAXPHYADDR and whether the PSE-36 mechanism is supported.

## PAGING

- Bits 11:2 are bits 31:22 of the linear address.
- Bits 1:0 are 0.

Because a PDE is identified using bits 31:22 of the linear address, it controls access to a 4-Mbyte region of the linear-address space. Use of the PDE depends on CR4.PSE and the PDE's PS flag (bit 7):

- If CR4.PSE = 1 and the PDE's PS flag is 1, the PDE maps a 4-MByte page (see Table 4-4). The final physical address is computed as follows:
  - Bits 39:32 are bits 20:13 of the PDE.
  - Bits 31:22 are bits 31:22 of the PDE.<sup>1</sup>
  - Bits 21:0 are from the original linear address.
- If CR4.PSE = 0 or the PDE's PS flag is 0, a 4-KByte naturally aligned page table is located at the physical address specified in bits 31:12 of the PDE (see Table 4-5). A page table comprises 1024 32-bit entries (PTEs). A PTE is selected using the physical address defined as follows:
  - Bits 39:32 are all 0.
  - Bits 31:12 are from the PDE.
  - Bits 11:2 are bits 21:12 of the linear address.
  - Bits 1:0 are 0.
- Because a PTE is identified using bits 31:12 of the linear address, every PTE maps a 4-KByte page (see Table 4-6). The final physical address is computed as follows:
  - Bits 39:32 are all 0.
  - Bits 31:12 are from the PTE.
  - Bits 11:0 are from the original linear address.

If a paging-structure entry's P flag (bit 0) is 0 or if the entry sets any reserved bit, the entry is used neither to reference another paging-structure entry nor to map a page. There is no translation for a linear address whose translation would use such a paging-structure entry; a reference to such a linear address causes a page-fault exception (see Section 4.7).

With 32-bit paging, there are reserved bits only if CR4.PSE = 1:

- If the P flag and the PS flag (bit 7) of a PDE are both 1, the bits reserved depend on MAXPHYADDR, and whether the PSE-36 mechanism is supported:<sup>2</sup>
  - If the PSE-36 mechanism is not supported, bits 21:13 are reserved.
  - If the PSE-36 mechanism is supported, bits 21:(M-19) are reserved, where M is the minimum of 40 and MAXPHYADDR.
- If the PAT is not supported:<sup>3</sup>
  - If the P flag of a PTE is 1, bit 7 is reserved.
  - If the P flag and the PS flag of a PDE are both 1, bit 12 is reserved.

(If CR4.PSE = 0, no bits are reserved with 32-bit paging.)

A reference using a linear address that is successfully translated to a physical address is performed only if allowed by the access rights of the translation; see Section 4.6.

---

1. The upper bits in the final physical address do not all come from corresponding positions in the PDE; the physical-address bits in the PDE are not all contiguous.

2. See Section 4.1.4 for how to determine MAXPHYADDR and whether the PSE-36 mechanism is supported.

3. See Section 4.1.4 for how to determine whether the PAT is supported.



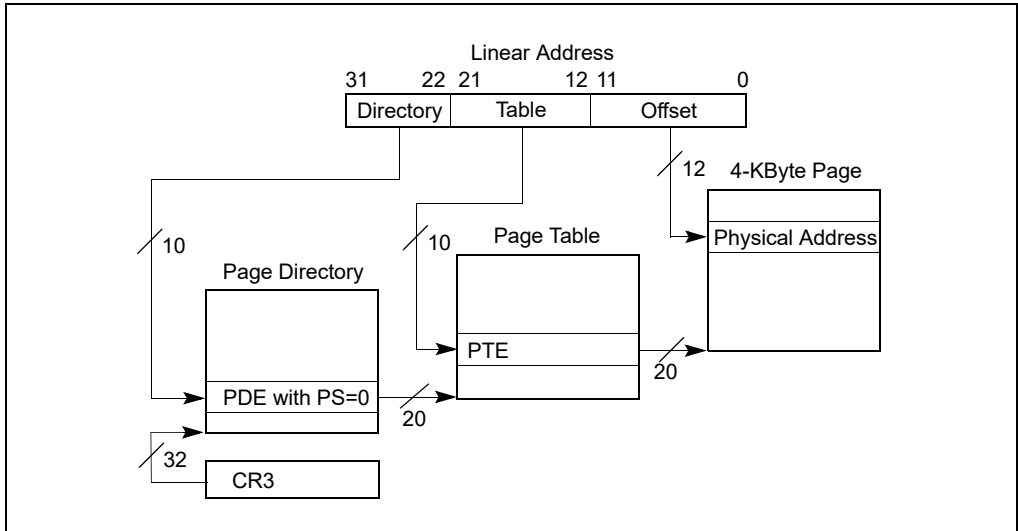


Figure 4-2. Linear-Address Translation to a 4-KByte Page using 32-Bit Paging

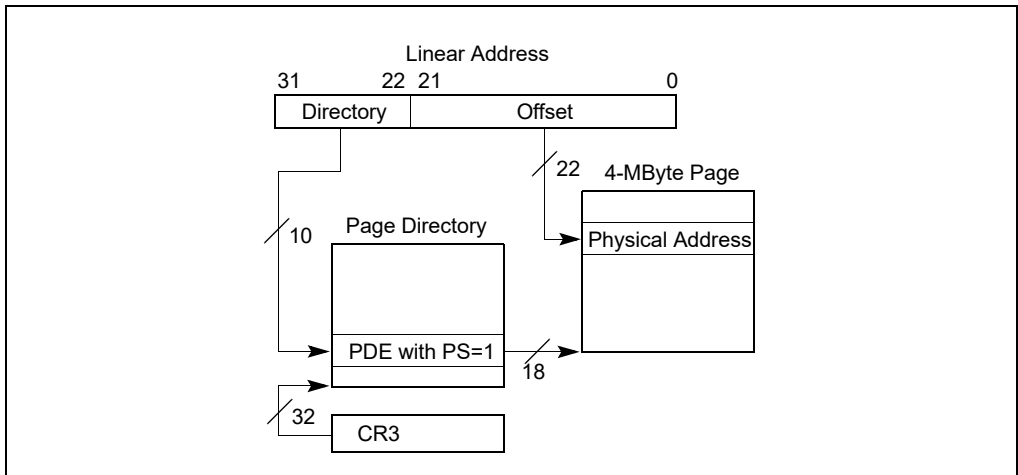


Figure 4-3. Linear-Address Translation to a 4-MByte Page using 32-Bit Paging

Figure 4-4 gives a summary of the formats of CR3 and the paging-structure entries with 32-bit paging. For the paging structure entries, it identifies separately the format of entries that map pages, those that reference other paging structures, and those that do neither because they are “not present”; bit 0 (P) and bit 7 (PS) are highlighted because they determine how such an entry is used.

|   |    |    |    |    |    |                      |    |    |                                    |    |    |         |         |    |          |    |    |          |                  |         |     |          |               |     |          |                 |               |   |   |   |   |  |
|---|----|----|----|----|----|----------------------|----|----|------------------------------------|----|----|---------|---------|----|----------|----|----|----------|------------------|---------|-----|----------|---------------|-----|----------|-----------------|---------------|---|---|---|---|--|
| 31                                      | 30 | 29 | 28 | 27 | 26 | 25                   | 24 | 23 | 22                                 | 21 | 20 | 19      | 18      | 17 | 16       | 15 | 14 | 13       | 12               | 11      | 10  | 9        | 8             | 7   | 6        | 5               | 4             | 3 | 2 | 1 | 0 |  |
| Address of page directory <sup>1</sup>  |    |    |    |    |    |                      |    |    |                                    |    |    | Ignored |         |    |          |    |    | PCD      | PWT              | Ignored |     |          | CR3           |     |          |                 |               |   |   |   |   |  |
| Bits 31:22 of address of 4MB page frame |    |    |    |    |    | Reserved (must be 0) |    |    | Bits 39:32 of address <sup>2</sup> |    |    | PAT     | Ignored | G  | <u>1</u> | D  | A  | PCD      | PWT              | U/S     | R/W | <u>1</u> | PDE: 4MB page |     |          |                 |               |   |   |   |   |  |
| Address of page table                   |    |    |    |    |    |                      |    |    |                                    |    |    | Ignored |         |    |          |    |    | <u>0</u> | Ign              | A       | PCD | PWT      | U/S           | R/W | <u>1</u> | PDE: page table |               |   |   |   |   |  |
| Ignored                                 |    |    |    |    |    |                      |    |    |                                    |    |    |         |         |    |          |    |    | <u>0</u> | PDE: not present |         |     |          |               |     |          |                 |               |   |   |   |   |  |
| Address of 4KB page frame               |    |    |    |    |    |                      |    |    |                                    |    |    | Ignored |         |    |          |    |    | G        | PAT              | D       | A   | PCD      | PWT           | U/S | R/W      | <u>1</u>        | PTE: 4KB page |   |   |   |   |  |
| Ignored                                 |    |    |    |    |    |                      |    |    |                                    |    |    |         |         |    |          |    |    | <u>0</u> | PTE: not present |         |     |          |               |     |          |                 |               |   |   |   |   |  |

Figure 4-4. Formats of CR3 and Paging-Structure Entries with 32-Bit Paging

NOTES:

1. CR3 has 64 bits on processors supporting the Intel-64 architecture. These bits are ignored with 32-bit paging.
2. This example illustrates a processor in which MAXPHYADDR is 36. If this value is larger or smaller, the number of bits reserved in positions 20:13 of a PDE mapping a 4-MByte page will change.

Table 4-3. Use of CR3 with 32-Bit Paging

| Bit Position(s) | Contents  |
|-----------------|---|
| 2:0             | Ignored   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page directory during linear-address translation (see Section 4.9) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page directory during linear-address translation (see Section 4.9) |
| 11:5            | Ignored   |
| 31:12           | Physical address of the 4-KByte aligned page directory used for linear-address translation  |
| 63:32           | Ignored (these bits exist only on processors supporting the Intel-64 architecture)  |

**Table 4-4. Format of a 32-Bit Page-Directory Entry that Maps a 4-MByte Page**

| Bit Position(s) | Contents  |
|-----------------|---|
| 0 (P)           | Present; must be 1 to map a 4-MByte page  |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 4-MByte page referenced by this entry (see Section 4.6)  |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 4-MByte page referenced by this entry (see Section 4.6)  |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 4-MByte page referenced by this entry (see Section 4.9)  |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 4-MByte page referenced by this entry (see Section 4.9)  |
| 5 (A)           | Accessed; indicates whether software has accessed the 4-MByte page referenced by this entry (see Section 4.8)   |
| 6 (D)           | Dirty; indicates whether software has written to the 4-MByte page referenced by this entry (see Section 4.8)  |
| 7 (PS)          | Page size; must be 1 (otherwise, this entry references a page table; see Table 4-5)   |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise  |
| 11:9            | Ignored   |
| 12 (PAT)        | If the PAT is supported, indirectly determines the memory type used to access the 4-MByte page referenced by this entry (see Section 4.9.2); otherwise, reserved (must be 0) <sup>1</sup> |
| (M-20):13       | Bits (M-1):32 of physical address of the 4-MByte page referenced by this entry <sup>2</sup>   |
| 21:(M-19)       | Reserved (must be 0)  |
| 31:22           | Bits 31:22 of physical address of the 4-MByte page referenced by this entry   |

**NOTES:**

1. See Section 4.1.4 for how to determine whether the PAT is supported.
2. If the PSE-36 mechanism is not supported, M is 32, and this row does not apply. If the PSE-36 mechanism is supported, M is the minimum of 40 and MAXPHYADDR (this row does not apply if MAXPHYADDR = 32). See Section 4.1.4 for how to determine MAXPHYADDR and whether the PSE-36 mechanism is supported.

**Table 4-5. Format of a 32-Bit Page-Directory Entry that References a Page Table**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page table   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 4-MByte region controlled by this entry (see Section 4.6)                             |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 4-MByte region controlled by this entry (see Section 4.6)               |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9) |
| 5 (A)           | Accessed; indicates whether this entry has been used for linear-address translation (see Section 4.8)                                    |
| 6               | Ignored  |
| 7 (PS)          | If CR4.PSE = 1, must be 0 (otherwise, this entry maps a 4-MByte page; see Table 4-4); otherwise, ignored                                 |
| 11:8            | Ignored  |
| 31:12           | Physical address of 4-KByte aligned page table referenced by this entry  |

**Table 4-6. Format of a 32-Bit Page-Table Entry that Maps a 4-KByte Page**

| Bit Position(s) | Contents  |
|-----------------|---|
| 0 (P)           | Present; must be 1 to map a 4-KByte page  |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 4-KByte page referenced by this entry (see Section 4.6)  |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 4-KByte page referenced by this entry (see Section 4.6)  |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9)  |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9)  |
| 5 (A)           | Accessed; indicates whether software has accessed the 4-KByte page referenced by this entry (see Section 4.8)   |
| 6 (D)           | Dirty; indicates whether software has written to the 4-KByte page referenced by this entry (see Section 4.8)  |
| 7 (PAT)         | If the PAT is supported, indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9.2); otherwise, reserved (must be 0) <sup>1</sup> |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise  |
| 11:9            | Ignored   |
| 31:12           | Physical address of the 4-KByte page referenced by this entry   |

**NOTES:**

1. See Section 4.1.4 for how to determine whether the PAT is supported.

## 4.4 PAE PAGING

A logical processor uses PAE paging if  $CR0.PG = 1$ ,  $CR4.PAE = 1$ , and  $IA32\_EFER.LME = 0$ . PAE paging translates 32-bit linear addresses to 52-bit physical addresses.<sup>1</sup> Although 52 bits corresponds to 4 PBytes, linear addresses are limited to 32 bits; at most 4 GBytes of linear-address space may be accessed at any given time.

With PAE paging, a logical processor maintains a set of four (4) PDPTE registers, which are loaded from an address in CR3. Linear address are translated using 4 hierarchies of in-memory paging structures, each located using one of the PDPTE registers. (This is different from the other paging modes, in which there is one hierarchy referenced by CR3.)

Section 4.4.1 discusses the PDPTE registers. Section 4.4.2 describes linear-address translation with PAE paging.

### 4.4.1 PDPTE Registers

When PAE paging is used, CR3 references the base of a 32-Byte page-directory-pointer table. Table 4-7 illustrates how CR3 is used with PAE paging.

**Table 4-7. Use of CR3 with PAE Paging**

| Bit Position(s) | Contents   |
|-----------------|--|
| 4:0             | Ignored  |
| 31:5            | Physical address of the 32-Byte aligned page-directory-pointer table used for linear-address translation |
| 63:32           | Ignored (these bits exist only on processors supporting the Intel-64 architecture)                       |

The page-directory-pointer-table comprises four (4) 64-bit entries called PDPTes. Each PDPTE controls access to a 1-GByte region of the linear-address space. Corresponding to the PDPTes, the logical processor maintains a set of four (4) internal, non-architectural PDPTE registers, called PDPTE0, PDPTE1, PDPTE2, and PDPTE3. The logical processor loads these registers from the PDPTes in memory as part of certain operations:

- If PAE paging would be in use following an execution of MOV to CR0 or MOV to CR4 (see Section 4.1.1) and the instruction is modifying any of CR0.CD, CR0.NW, CR0.PG, CR4.PAE, CR4.PGE, CR4.PSE, or CR4.SMEP; then the PDPTes are loaded from the address in CR3.
- If MOV to CR3 is executed while the logical processor is using PAE paging, the PDPTes are loaded from the address being loaded into CR3.
- If PAE paging is in use and a task switch changes the value of CR3, the PDPTes are loaded from the address in the new CR3 value.
- Certain VMX transitions load the PDPTE registers. See Section 4.11.1.

Table 4-8 gives the format of a PDPTE. If any of the PDPTes sets both the P flag (bit 0) and any reserved bit, the MOV to CR instruction causes a general-protection exception (#GP(0)) and the PDPTes are not loaded.<sup>2</sup> As shown in Table 4-8, bits 2:1, 8:5, and 63:MAXPHYADDR are reserved in the PDPTes.

1. If  $MAXPHYADDR < 52$ , bits in the range 51:MAXPHYADDR will be 0 in any physical address used by PAE paging. (The corresponding bits are reserved in the paging-structure entries.) See Section 4.1.4 for how to determine MAXPHYADDR.

2. On some processors, reserved bits are checked even in PDPTes in which the P flag (bit 0) is 0.

**Table 4-8. Format of a PAE Page-Directory-Pointer-Table Entry (PDPTE)**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page directory   |
| 2:1             | Reserved (must be 0)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page directory referenced by this entry (see Section 4.9) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page directory referenced by this entry (see Section 4.9) |
| 8:5             | Reserved (must be 0)   |
| 11:9            | Ignored  |
| (M-1):12        | Physical address of 4-KByte aligned page directory referenced by this entry <sup>1</sup>   |
| 63:M            | Reserved (must be 0)   |

**NOTES:**

1. M is an abbreviation for MAXPHYADDR, which is at most 52; see Section 4.1.4.

#### 4.4.2 Linear-Address Translation with PAE Paging

PAE paging may map linear addresses to either 4-KByte pages or 2-MByte pages. Figure 4-5 illustrates the translation process when it produces a 4-KByte page; Figure 4-6 covers the case of a 2-MByte page. The following items describe the PAE paging process in more detail as well as how the page size is determined:

- Bits 31:30 of the linear address select a PDPTE register (see Section 4.4.1); this is PDPTE $_i$ , where  $i$  is the value of bits 31:30.<sup>1</sup> Because a PDPTE register is identified using bits 31:30 of the linear address, it controls access to a 1-GByte region of the linear-address space. If the P flag (bit 0) of PDPTE $_i$  is 0, the processor ignores bits 63:1, and there is no mapping for the 1-GByte region controlled by PDPTE $_i$ . A reference using a linear address in this region causes a page-fault exception (see Section 4.7).
- If the P flag of PDPTE $_i$  is 1, 4-KByte naturally aligned page directory is located at the physical address specified in bits 51:12 of PDPTE $_i$  (see Table 4-8 in Section 4.4.1). A page directory comprises 512 64-bit entries (PDEs). A PDE is selected using the physical address defined as follows:
  - Bits 51:12 are from PDPTE $_i$ .
  - Bits 11:3 are bits 29:21 of the linear address.
  - Bits 2:0 are 0.

Because a PDE is identified using bits 31:21 of the linear address, it controls access to a 2-Mbyte region of the linear-address space. Use of the PDE depends on its PS flag (bit 7):

- If the PDE's PS flag is 1, the PDE maps a 2-MByte page (see Table 4-9). The final physical address is computed as follows:
  - Bits 51:21 are from the PDE.
  - Bits 20:0 are from the original linear address.
- If the PDE's PS flag is 0, a 4-KByte naturally aligned page table is located at the physical address specified in bits 51:12 of the PDE (see Table 4-10). A page table comprises 512 64-bit entries (PTEs). A PTE is selected using the physical address defined as follows:
  - Bits 51:12 are from the PDE.

1. With PAE paging, the processor does not use CR3 when translating a linear address (as it does in the other paging modes). It does not access the PDPTes in the page-directory-pointer table during linear-address translation.

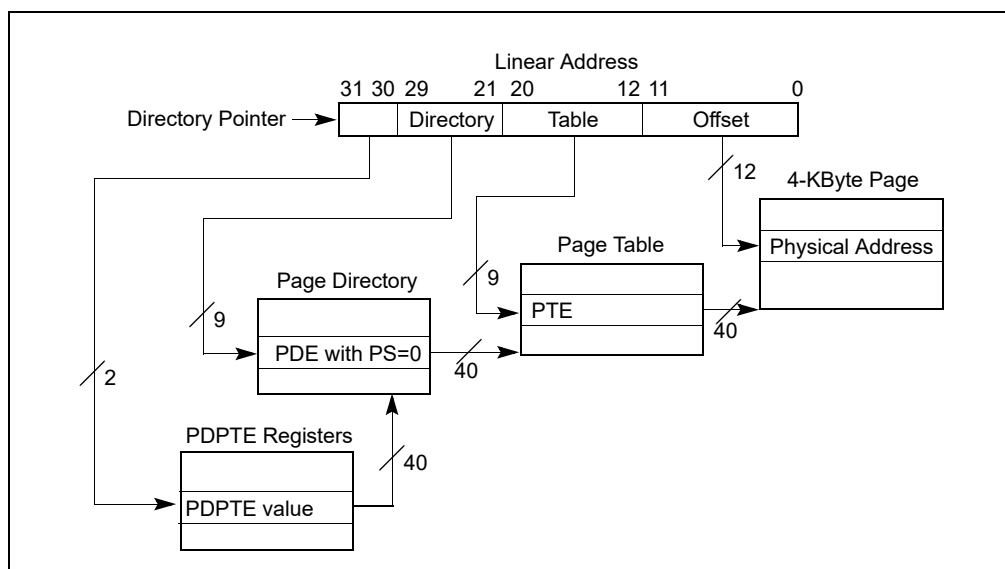
- Bits 11:3 are bits 20:12 of the linear address.
- Bits 2:0 are 0.
- Because a PTE is identified using bits 31:12 of the linear address, every PTE maps a 4-KByte page (see Table 4-11). The final physical address is computed as follows:
  - Bits 51:12 are from the PTE.
  - Bits 11:0 are from the original linear address.

If the P flag (bit 0) of a PDE or a PTE is 0 or if a PDE or a PTE sets any reserved bit, the entry is used neither to reference another paging-structure entry nor to map a page. There is no translation for a linear address whose translation would use such a paging-structure entry; a reference to such a linear address causes a page-fault exception (see Section 4.7).

The following bits are reserved with PAE paging:

- If the P flag (bit 0) of a PDE or a PTE is 1, bits 62:MAXPHYADDR are reserved.
- If the P flag and the PS flag (bit 7) of a PDE are both 1, bits 20:13 are reserved.
- If IA32\_EFER.NXE = 0 and the P flag of a PDE or a PTE is 1, the XD flag (bit 63) is reserved.
- If the PAT is not supported:<sup>1</sup>
  - If the P flag of a PTE is 1, bit 7 is reserved.
  - If the P flag and the PS flag of a PDE are both 1, bit 12 is reserved.

A reference using a linear address that is successfully translated to a physical address is performed only if allowed by the access rights of the translation; see Section 4.6.



**Figure 4-5. Linear-Address Translation to a 4-KByte Page using PAE Paging**

1. See Section 4.1.4 for how to determine whether the PAT is supported.

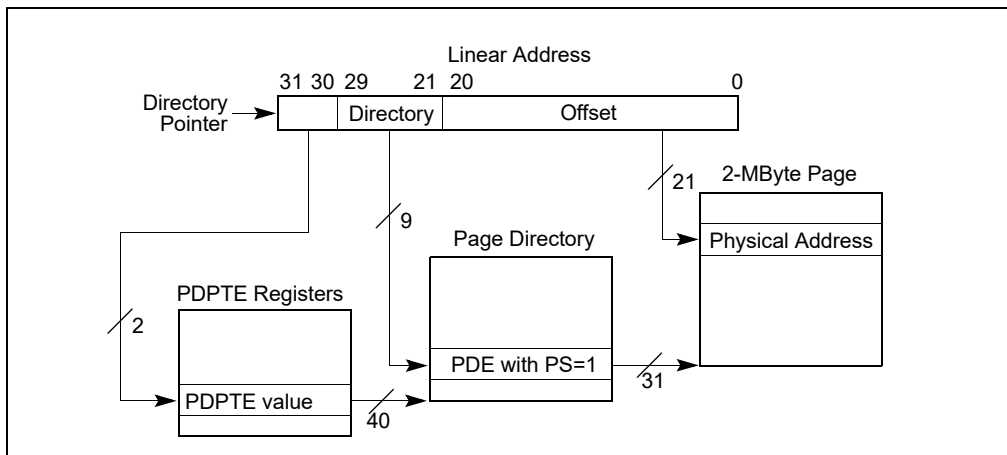


Figure 4-6. Linear-Address Translation to a 2-MByte Page using PAE Paging

Table 4-9. Format of a PAE Page-Directory Entry that Maps a 2-MByte Page

| Bit Position(s) | Contents  |
|-----------------|---|
| 0 (P)           | Present; must be 1 to map a 2-MByte page  |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 2-MByte page referenced by this entry (see Section 4.6)  |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 2-MByte page referenced by this entry (see Section 4.6)  |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9)  |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9)  |
| 5 (A)           | Accessed; indicates whether software has accessed the 2-MByte page referenced by this entry (see Section 4.8)   |
| 6 (D)           | Dirty; indicates whether software has written to the 2-MByte page referenced by this entry (see Section 4.8)  |
| 7 (PS)          | Page size; must be 1 (otherwise, this entry references a page table; see Table 4-10)  |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise  |
| 11:9            | Ignored   |
| 12 (PAT)        | If the PAT is supported, indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9.2); otherwise, reserved (must be 0) <sup>1</sup> |
| 20:13           | Reserved (must be 0)  |
| (M-1):21        | Physical address of the 2-MByte page referenced by this entry   |
| 62:M            | Reserved (must be 0)  |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 2-MByte page controlled by this entry; see Section 4.6); otherwise, reserved (must be 0)        |

**NOTES:**

1. See Section 4.1.4 for how to determine whether the PAT is supported.



**Table 4-10. Format of a PAE Page-Directory Entry that References a Page Table**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page table   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 2-MByte region controlled by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 2-MByte region controlled by this entry (see Section 4.6)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9)   |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9)   |
| 5 (A)           | Accessed; indicates whether this entry has been used for linear-address translation (see Section 4.8)  |
| 6               | Ignored  |
| 7 (PS)          | Page size; must be 0 (otherwise, this entry maps a 2-MByte page; see Table 4-9)  |
| 11:8            | Ignored  |
| (M-1):12        | Physical address of 4-KByte aligned page table referenced by this entry  |
| 62:M            | Reserved (must be 0)   |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 2-MByte region controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**Table 4-11. Format of a PAE Page-Table Entry that Maps a 4-KByte Page**

| Bit Position(s) | Contents  |
|-----------------|---|
| 0 (P)           | Present; must be 1 to map a 4-KByte page  |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 4-KByte page referenced by this entry (see Section 4.6)  |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 4-KByte page referenced by this entry (see Section 4.6)  |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9)  |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9)  |
| 5 (A)           | Accessed; indicates whether software has accessed the 4-KByte page referenced by this entry (see Section 4.8)   |
| 6 (D)           | Dirty; indicates whether software has written to the 4-KByte page referenced by this entry (see Section 4.8)  |
| 7 (PAT)         | If the PAT is supported, indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9.2); otherwise, reserved (must be 0) <sup>1</sup> |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise  |



## 4.5 4-LEVEL PAGING

A logical processor uses 4-level paging if CR0.PG = 1, CR4.PAE = 1, and IA32\_EFER.LME = 1. With 4-level paging, linear addresses are translated using a hierarchy of in-memory paging structures located using the contents of CR3. 4-level paging translates 48-bit linear addresses to 52-bit physical addresses.<sup>1</sup> Although 52 bits corresponds to 4 PBytes, linear addresses are limited to 48 bits; at most 256 TBytes of linear-address space may be accessed at any given time.

4-level paging uses a hierarchy of paging structures to produce a translation for a linear address. CR3 is used to locate the first paging-structure, the PML4 table. Use of CR3 with 4-level paging depends on whether process-context identifiers (PCIDs) have been enabled by setting CR4.PCIDE:

- Table 4-12 illustrates how CR3 is used with 4-level paging if CR4.PCIDE = 0.

**Table 4-12. Use of CR3 with 4-Level Paging and CR4.PCIDE = 0**

| Bit Position(s) | Contents  |
|-----------------|---|
| 2:0             | Ignored   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the PML4 table during linear-address translation (see Section 4.9.2) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the PML4 table during linear-address translation (see Section 4.9.2) |
| 11:5            | Ignored   |
| M-1:12          | Physical address of the 4-KByte aligned PML4 table used for linear-address translation <sup>1</sup>   |
| 63:M            | Reserved (must be 0)  |

### NOTES:

1. M is an abbreviation for MAXPHYADDR, which is at most 52; see Section 4.1.4.

- Table 4-13 illustrates how CR3 is used with 4-level paging if CR4.PCIDE = 1.

**Table 4-13. Use of CR3 with 4-Level Paging and CR4.PCIDE = 1**

| Bit Position(s) | Contents  |
|-----------------|---|
| 11:0            | PCID (see Section 4.10.1) <sup>1</sup>  |
| M-1:12          | Physical address of the 4-KByte aligned PML4 table used for linear-address translation <sup>2</sup> |
| 63:M            | Reserved (must be 0) <sup>3</sup>   |

### NOTES:

1. Section 4.9.2 explains how the processor determines the memory type used to access the PML4 table during linear-address translation with CR4.PCIDE = 1.

2. M is an abbreviation for MAXPHYADDR, which is at most 52; see Section 4.1.4.

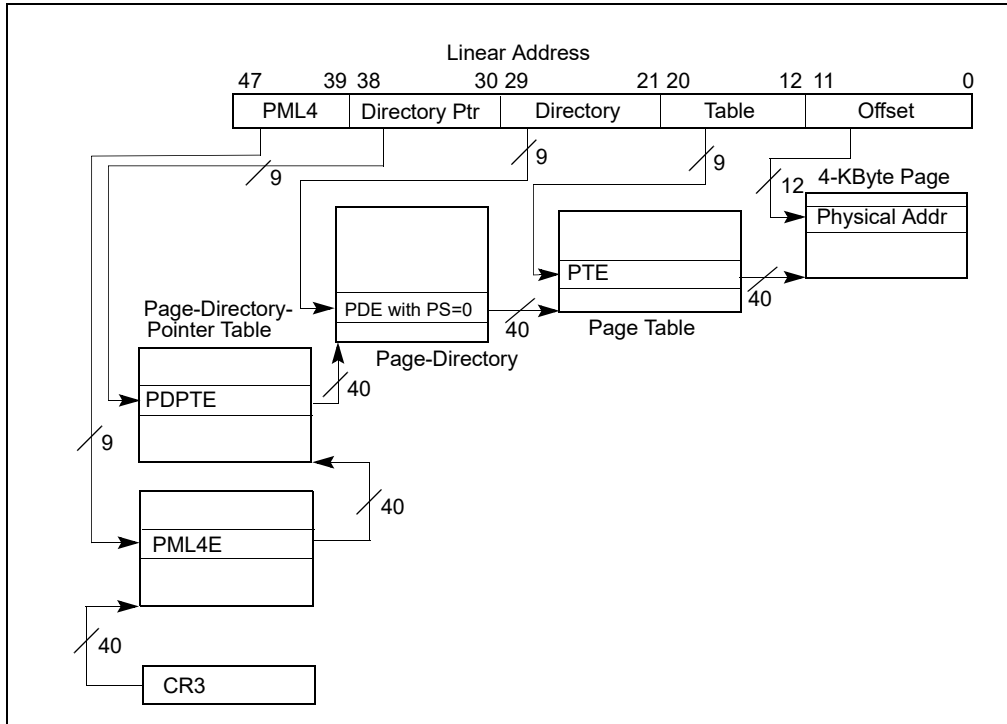
3. See Section 4.10.4.1 for use of bit 63 of the source operand of the MOV to CR3 instruction.

After software modifies the value of CR4.PCIDE, the logical processor immediately begins using CR3 as specified for the new value. For example, if software changes CR4.PCIDE from 1 to 0, the current PCID immediately changes

1. If MAXPHYADDR < 52, bits in the range 51:MAXPHYADDR will be 0 in any physical address used by 4-level paging. (The corresponding bits are reserved in the paging-structure entries.) See Section 4.1.4 for how to determine MAXPHYADDR.

## PAGING

from CR3[11:0] to 000H (see also Section 4.10.4.1). In addition, the logical processor subsequently determines the memory type used to access the PML4 table using CR3.PWT and CR3.PCD, which had been bits 4:3 of the PCID. 4-level paging may map linear addresses to 4-KByte pages, 2-MByte pages, or 1-GByte pages.<sup>1</sup> Figure 4-8 illustrates the translation process when it produces a 4-KByte page; Figure 4-9 covers the case of a 2-MByte page, and Figure 4-10 the case of a 1-GByte page.



**Figure 4-8. Linear-Address Translation to a 4-KByte Page using 4-Level Paging**

1. Not all processors support 1-GByte pages; see Section 4.1.4.

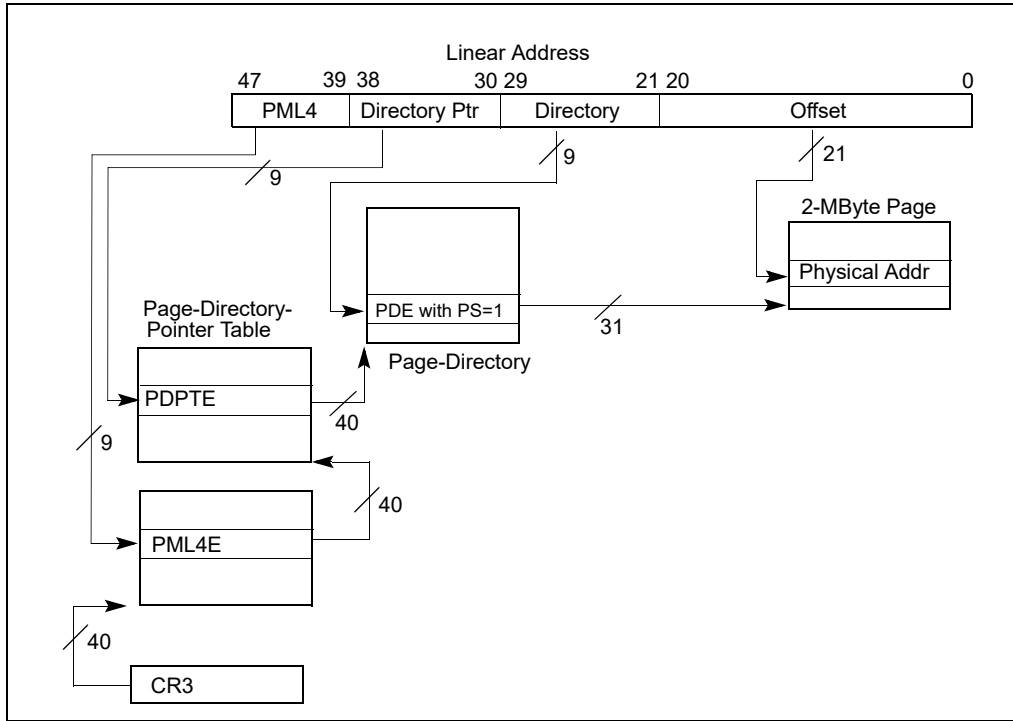


Figure 4-9. Linear-Address Translation to a 2-MByte Page using 4-Level Paging

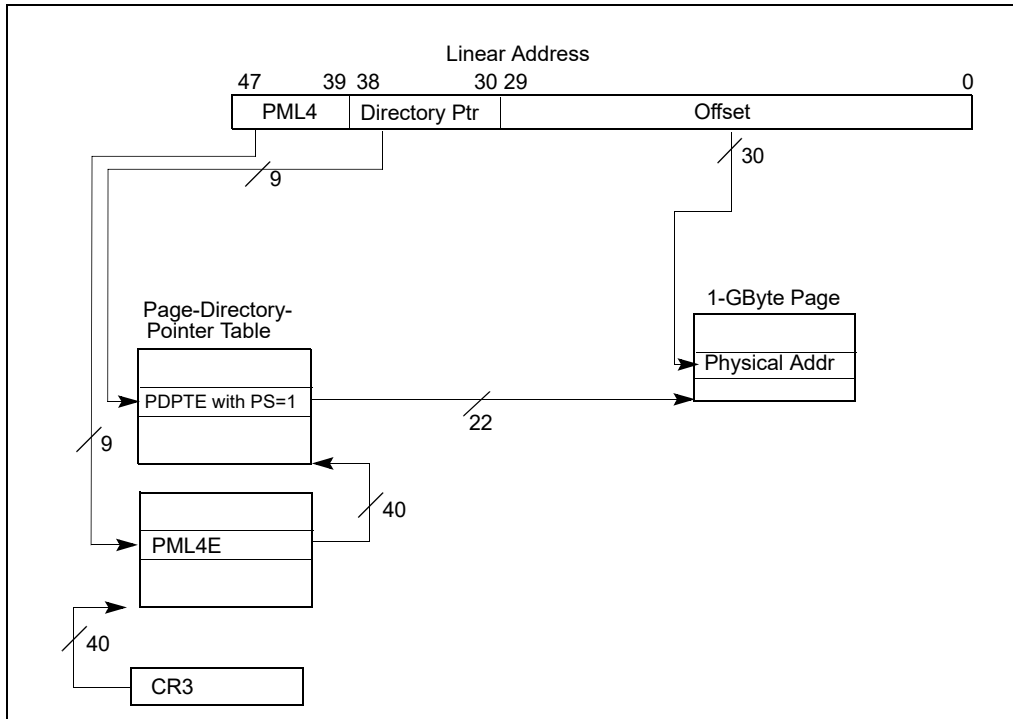


Figure 4-10. Linear-Address Translation to a 1-GByte Page using 4-Level Paging

If CR4.PKE = 1, 4-level paging associates with each linear address a **protection key**. Section 4.6 explains how the processor uses the protection key in its determination of the access rights of each linear address.

The following items describe the 4-level paging process in more detail as well as how the page size and protection key are determined.

- A 4-KByte naturally aligned PML4 table is located at the physical address specified in bits 51:12 of CR3 (see Table 4-12). A PML4 table comprises 512 64-bit entries (PML4Es). A PML4E is selected using the physical address defined as follows:
  - Bits 51:12 are from CR3.
  - Bits 11:3 are bits 47:39 of the linear address.
  - Bits 2:0 are all 0.

Because a PML4E is identified using bits 47:39 of the linear address, it controls access to a 512-GByte region of the linear-address space.
- A 4-KByte naturally aligned page-directory-pointer table is located at the physical address specified in bits 51:12 of the PML4E (see Table 4-14). A page-directory-pointer table comprises 512 64-bit entries (PDPTEs). A PDPTE is selected using the physical address defined as follows:
  - Bits 51:12 are from the PML4E.
  - Bits 11:3 are bits 38:30 of the linear address.
  - Bits 2:0 are all 0.

Because a PDPTE is identified using bits 47:30 of the linear address, it controls access to a 1-GByte region of the linear-address space. Use of the PDPTE depends on its PS flag (bit 7):<sup>1</sup>

- If the PDPTE's PS flag is 1, the PDPTE maps a 1-GByte page (see Table 4-15). The final physical address is computed as follows:
  - Bits 51:30 are from the PDPTE.
  - Bits 29:0 are from the original linear address.

If CR4.PKE = 1, the linear address's protection key is the value of bits 62:59 of the PDPTE.
- If the PDPTE's PS flag is 0, a 4-KByte naturally aligned page directory is located at the physical address specified in bits 51:12 of the PDPTE (see Table 4-16). A page directory comprises 512 64-bit entries (PDEs). A PDE is selected using the physical address defined as follows:
  - Bits 51:12 are from the PDPTE.
  - Bits 11:3 are bits 29:21 of the linear address.
  - Bits 2:0 are all 0.

Because a PDE is identified using bits 47:21 of the linear address, it controls access to a 2-MByte region of the linear-address space. Use of the PDE depends on its PS flag:

- If the PDE's PS flag is 1, the PDE maps a 2-MByte page (see Table 4-17). The final physical address is computed as follows:
  - Bits 51:21 are from the PDE.
  - Bits 20:0 are from the original linear address.

If CR4.PKE = 1, the linear address's protection key is the value of bits 62:59 of the PDE.
- If the PDE's PS flag is 0, a 4-KByte naturally aligned page table is located at the physical address specified in bits 51:12 of the PDE (see Table 4-18). A page table comprises 512 64-bit entries (PTEs). A PTE is selected using the physical address defined as follows:
  - Bits 51:12 are from the PDE.
  - Bits 11:3 are bits 20:12 of the linear address.
  - Bits 2:0 are all 0.

---

1. The PS flag of a PDPTE is reserved and must be 0 (if the P flag is 1) if 1-GByte pages are not supported. See Section 4.1.4 for how to determine whether 1-GByte pages are supported.

- Because a PTE is identified using bits 47:12 of the linear address, every PTE maps a 4-KByte page (see Table 4-19). The final physical address is computed as follows:
  - Bits 51:12 are from the PTE.
  - Bits 11:0 are from the original linear address.

If CR4.PKE = 1, the linear address's protection key is the value of bits 62:59 of the PTE.

If a paging-structure entry's P flag (bit 0) is 0 or if the entry sets any reserved bit, the entry is used neither to reference another paging-structure entry nor to map a page. There is no translation for a linear address whose translation would use such a paging-structure entry; a reference to such a linear address causes a page-fault exception (see Section 4.7).

The following bits are reserved with 4-level paging:

- If the P flag of a paging-structure entry is 1, bits 51:MAXPHYADDR are reserved.
- If the P flag of a PML4E is 1, the PS flag is reserved.
- If 1-GByte pages are not supported and the P flag of a PDPTE is 1, the PS flag is reserved.<sup>1</sup>
- If the P flag and the PS flag of a PDPTE are both 1, bits 29:13 are reserved.
- If the P flag and the PS flag of a PDE are both 1, bits 20:13 are reserved.
- If IA32\_EFER.NXE = 0 and the P flag of a paging-structure entry is 1, the XD flag (bit 63) is reserved.

A reference using a linear address that is successfully translated to a physical address is performed only if allowed by the access rights of the translation; see Section 4.6.

Figure 4-11 gives a summary of the formats of CR3 and the 4-level paging-structure entries. For the paging structure entries, it identifies separately the format of entries that map pages, those that reference other paging structures, and those that do neither because they are "not present"; bit 0 (P) and bit 7 (PS) are highlighted because they determine how a paging-structure entry is used.

**Table 4-14. Format of a 4-Level PML4 Entry (PML4E) that References a Page-Directory-Pointer Table**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page-directory-pointer table   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 512-GByte region controlled by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 512-GByte region controlled by this entry (see Section 4.6)                                 |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page-directory-pointer table referenced by this entry (see Section 4.9.2) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page-directory-pointer table referenced by this entry (see Section 4.9.2) |
| 5 (A)           | Accessed; indicates whether this entry has been used for linear-address translation (see Section 4.8)  |
| 6               | Ignored  |
| 7 (PS)          | Reserved (must be 0)   |
| 11:8            | Ignored  |

1. See Section 4.1.4 for how to determine whether 1-GByte pages are supported.

**Table 4-14. Format of a 4-Level PML4 Entry (PML4E) that References a Page-Directory-Pointer Table (Contd.)**

| Bit Position(s) | Contents   |
|-----------------|--|
| M-1:12          | Physical address of 4-KByte aligned page-directory-pointer table referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 62:52           | Ignored  |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 512-GByte region controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**Table 4-15. Format of a 4-Level Page-Directory-Pointer-Table Entry (PDPTE) that Maps a 1-GByte Page**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to map a 1-GByte page   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 1-GByte page referenced by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 1-GByte page referenced by this entry (see Section 4.6)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 1-GByte page referenced by this entry (see Section 4.9.2)                                       |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 1-GByte page referenced by this entry (see Section 4.9.2)                                       |
| 5 (A)           | Accessed; indicates whether software has accessed the 1-GByte page referenced by this entry (see Section 4.8)  |
| 6 (D)           | Dirty; indicates whether software has written to the 1-GByte page referenced by this entry (see Section 4.8)   |
| 7 (PS)          | Page size; must be 1 (otherwise, this entry references a page directory; see Table 4-16)   |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise   |
| 11:9            | Ignored  |
| 12 (PAT)        | Indirectly determines the memory type used to access the 1-GByte page referenced by this entry (see Section 4.9.2) <sup>1</sup>  |
| 29:13           | Reserved (must be 0)   |
| (M-1):30        | Physical address of the 1-GByte page referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 58:52           | Ignored  |
| 62:59           | Protection key; if CR4.PKE = 1, determines the protection key of the page (see Section 4.6.2); ignored otherwise   |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 1-GByte page controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**NOTES:**

1. The PAT is supported on all processors that support 4-level paging.



**Table 4-16. Format of a 4-Level Page-Directory-Pointer-Table Entry (PDPTE) that References a Page Directory**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page directory   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 1-GByte region controlled by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 1-GByte region controlled by this entry (see Section 4.6)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page directory referenced by this entry (see Section 4.9.2)                                       |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page directory referenced by this entry (see Section 4.9.2)                                       |
| 5 (A)           | Accessed; indicates whether this entry has been used for linear-address translation (see Section 4.8)  |
| 6               | Ignored  |
| 7 (PS)          | Page size; must be 0 (otherwise, this entry maps a 1-GByte page; see Table 4-15)   |
| 11:8            | Ignored  |
| (M-1):12        | Physical address of 4-KByte aligned page directory referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 62:52           | Ignored  |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 1-GByte region controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**Table 4-17. Format of a 4-Level Page-Directory Entry that Maps a 2-MByte Page**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to map a 2-MByte page   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 2-MByte page referenced by this entry (see Section 4.6)                                   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 2-MByte page referenced by this entry (see Section 4.6)                     |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9.2) |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9.2) |
| 5 (A)           | Accessed; indicates whether software has accessed the 2-MByte page referenced by this entry (see Section 4.8)                                |
| 6 (D)           | Dirty; indicates whether software has written to the 2-MByte page referenced by this entry (see Section 4.8)                                 |
| 7 (PS)          | Page size; must be 1 (otherwise, this entry references a page table; see Table 4-18)   |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise                                   |

**Table 4-17. Format of a 4-Level Page-Directory Entry that Maps a 2-MByte Page (Contd.)**

| Bit Position(s) | Contents   |
|-----------------|--|
| 11:9            | Ignored  |
| 12 (PAT)        | Indirectly determines the memory type used to access the 2-MByte page referenced by this entry (see Section 4.9.2)   |
| 20:13           | Reserved (must be 0)   |
| (M-1):21        | Physical address of the 2-MByte page referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 58:52           | Ignored  |
| 62:59           | Protection key; if CR4.PKE = 1, determines the protection key of the page (see Section 4.6.2); ignored otherwise   |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 2-MByte page controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**Table 4-18. Format of a 4-Level Page-Directory Entry that References a Page Table**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to reference a page table   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 2-MByte region controlled by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 2-MByte region controlled by this entry (see Section 4.6)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9.2)   |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the page table referenced by this entry (see Section 4.9.2)   |
| 5 (A)           | Accessed; indicates whether this entry has been used for linear-address translation (see Section 4.8)  |
| 6               | Ignored  |
| 7 (PS)          | Page size; must be 0 (otherwise, this entry maps a 2-MByte page; see Table 4-17)   |
| 11:8            | Ignored  |
| (M-1):12        | Physical address of 4-KByte aligned page table referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 62:52           | Ignored  |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 2-MByte region controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |

**Table 4-19. Format of a 4-Level Page-Table Entry that Maps a 4-KByte Page**

| Bit Position(s) | Contents   |
|-----------------|--|
| 0 (P)           | Present; must be 1 to map a 4-KByte page   |
| 1 (R/W)         | Read/write; if 0, writes may not be allowed to the 4-KByte page referenced by this entry (see Section 4.6)   |
| 2 (U/S)         | User/supervisor; if 0, user-mode accesses are not allowed to the 4-KByte page referenced by this entry (see Section 4.6)   |
| 3 (PWT)         | Page-level write-through; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9.2)                                       |
| 4 (PCD)         | Page-level cache disable; indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9.2)                                       |
| 5 (A)           | Accessed; indicates whether software has accessed the 4-KByte page referenced by this entry (see Section 4.8)  |
| 6 (D)           | Dirty; indicates whether software has written to the 4-KByte page referenced by this entry (see Section 4.8)   |
| 7 (PAT)         | Indirectly determines the memory type used to access the 4-KByte page referenced by this entry (see Section 4.9.2)   |
| 8 (G)           | Global; if CR4.PGE = 1, determines whether the translation is global (see Section 4.10); ignored otherwise   |
| 11:9            | Ignored  |
| (M-1):12        | Physical address of the 4-KByte page referenced by this entry  |
| 51:M            | Reserved (must be 0)   |
| 58:52           | Ignored  |
| 62:59           | Protection key; if CR4.PKE = 1, determines the protection key of the page (see Section 4.6.2); ignored otherwise   |
| 63 (XD)         | If IA32_EFER.NXE = 1, execute-disable (if 1, instruction fetches are not allowed from the 4-KByte page controlled by this entry; see Section 4.6); otherwise, reserved (must be 0) |



Section 4.6.1 describes how the processor determines the access rights for each linear address. Section 4.6.2 provides additional information about how protection keys contribute to access-rights determination. (They do so only with 4-level paging and only if CR4.PKE = 1.)

## 4.6.1 Determination of Access Rights

Every access to a linear address is either a **supervisor-mode** access or a **user-mode** access. For all instruction fetches and most data accesses, this distinction is determined by the current privilege level (CPL): accesses made while  $CPL < 3$  are supervisor-mode accesses, while accesses made while  $CPL = 3$  are user-mode accesses.

Some operations implicitly access system data structures with linear addresses; the resulting accesses to those data structures are supervisor-mode accesses regardless of CPL. Examples of such accesses include the following: accesses to the global descriptor table (GDT) or local descriptor table (LDT) to load a segment descriptor; accesses to the interrupt descriptor table (IDT) when delivering an interrupt or exception; and accesses to the task-state segment (TSS) as part of a task switch or change of CPL. All these accesses are called **implicit supervisor-mode** accesses regardless of CPL. Other accesses made while  $CPL < 3$  are called **explicit supervisor-mode** accesses.

Access rights are also controlled by the **mode** of a linear address as specified by the paging-structure entries controlling the translation of the linear address. If the U/S flag (bit 2) is 0 in at least one of the paging-structure entries, the address is a **supervisor-mode** address. Otherwise, the address is a **user-mode** address.

The following items detail how paging determines access rights:

- For supervisor-mode accesses:
    - Data may be read (implicitly or explicitly) from any supervisor-mode address.
    - Data reads from user-mode pages.  
Access rights depend on the value of CR4.SMAP:
      - If CR4.SMAP = 0, data may be read from any user-mode address with a protection key for which read access is permitted.
      - If CR4.SMAP = 1, access rights depend on the value of EFLAGS.AC and whether the access is implicit or explicit:
        - If EFLAGS.AC = 1 and the access is explicit, data may be read from any user-mode address with a protection key for which read access is permitted.
        - If EFLAGS.AC = 0 or the access is implicit, data may not be read from any user-mode address.
- Section 4.6.2 explains how protection keys are associated with user-mode addresses and the accesses that are permitted for each protection key.
- Data writes to supervisor-mode addresses.  
Access rights depend on the value of CR0.WP:
    - If CR0.WP = 0, data may be written to any supervisor-mode address.
    - If CR0.WP = 1, data may be written to any supervisor-mode address with a translation for which the R/W flag (bit 1) is 1 in every paging-structure entry controlling the translation; data may not be written to any supervisor-mode address with a translation for which the R/W flag is 0 in any paging-structure entry controlling the translation.
  - Data writes to user-mode addresses.  
Access rights depend on the value of CR0.WP:
    - If CR0.WP = 0, access rights depend on the value of CR4.SMAP:
      - If CR4.SMAP = 0, data may be written to any user-mode address with a protection key for which write access is permitted.
      - If CR4.SMAP = 1, access rights depend on the value of EFLAGS.AC and whether the access is implicit or explicit:
        - If EFLAGS.AC = 1 and the access is explicit, data may be written to any user-mode address with a protection key for which write access is permitted.
        - If EFLAGS.AC = 0 or the access is implicit, data may not be written to any user-mode address.

- If CR0.WP = 1, access rights depend on the value of CR4.SMAP:
  - If CR4.SMAP = 0, data may be written to any user-mode address with a translation for which the R/W flag is 1 in every paging-structure entry controlling the translation and with a protection key for which write access is permitted; data may not be written to any user-mode address with a translation for which the R/W flag is 0 in any paging-structure entry controlling the translation.
  - If CR4.SMAP = 1, access rights depend on the value of EFLAGS.AC and whether the access is implicit or explicit:
    - If EFLAGS.AC = 1 and the access is explicit, data may be written to any user-mode address with a translation for which the R/W flag is 1 in every paging-structure entry controlling the translation and with a protection key for which write access is permitted; data may not be written to any user-mode address with a translation for which the R/W flag is 0 in any paging-structure entry controlling the translation.
    - If EFLAGS.AC = 0 or the access is implicit, data may not be written to any user-mode address.

Section 4.6.2 explains how protection keys are associated with user-mode addresses and the accesses that are permitted for each protection key.

- Instruction fetches from supervisor-mode addresses.
  - For 32-bit paging or if IA32\_EFER.NXE = 0, instructions may be fetched from any supervisor-mode address.
  - For PAE paging or 4-level paging with IA32\_EFER.NXE = 1, instructions may be fetched from any supervisor-mode address with a translation for which the XD flag (bit 63) is 0 in every paging-structure entry controlling the translation; instructions may not be fetched from any supervisor-mode address with a translation for which the XD flag is 1 in any paging-structure entry controlling the translation.
- Instruction fetches from user-mode addresses.  
Access rights depend on the values of CR4.SMEP:
  - If CR4.SMEP = 0, access rights depend on the paging mode and the value of IA32\_EFER.NXE:
    - For 32-bit paging or if IA32\_EFER.NXE = 0, instructions may be fetched from any user-mode address.
    - For PAE paging or 4-level paging with IA32\_EFER.NXE = 1, instructions may be fetched from any user-mode address with a translation for which the XD flag is 0 in every paging-structure entry controlling the translation; instructions may not be fetched from any user-mode address with a translation for which the XD flag is 1 in any paging-structure entry controlling the translation.
  - If CR4.SMEP = 1, instructions may not be fetched from any user-mode address.
- For user-mode accesses:
  - Data reads.  
Access rights depend on the mode of the linear address:
    - Data may be read from any user-mode address with a protection key for which read access is permitted. Section 4.6.2 explains how protection keys are associated with user-mode addresses and the accesses that are permitted for each protection key.
    - Data may not be read from any supervisor-mode address.
  - Data writes.  
Access rights depend on the mode of the linear address:
    - Data may be written to any user-mode address with a translation for which the R/W flag is 1 in every paging-structure entry controlling the translation and with a protection key for which write access is permitted. Section 4.6.2 explains how protection keys are associated with user-mode addresses and the accesses that are permitted for each protection key.
    - Data may not be written to any supervisor-mode address.
  - Instruction fetches.  
Access rights depend on the mode of the linear address, the paging mode, and the value of IA32\_EFER.NXE:

- For 32-bit paging or if IA32\_EFER.NXE = 0, instructions may be fetched from any user-mode address.
- For PAE paging or 4-level paging with IA32\_EFER.NXE = 1, instructions may be fetched from any user-mode address with a translation for which the XD flag is 0 in every paging-structure entry controlling the translation.
- Instructions may not be fetched from any supervisor-mode address.

A processor may cache information from the paging-structure entries in TLBs and paging-structure caches (see Section 4.10). These structures may include information about access rights. The processor may enforce access rights based on the TLBs and paging-structure caches instead of on the paging structures in memory.

This fact implies that, if software modifies a paging-structure entry to change access rights, the processor might not use that change for a subsequent access to an affected linear address (see Section 4.10.4.3). See Section 4.10.4.2 for how software can ensure that the processor uses the modified access rights.

## 4.6.2 Protection Keys

The protection-key feature provides an additional mechanism by which 4-level paging controls access to user-mode addresses. When CR4.PKE = 1, every linear address is associated with the 4-bit **protection key** located in bits 62:59 of the paging-structure entry that mapped the page containing the linear address (see Section 4.5). The PKRU register determines, for each protection key, whether user-mode addresses with that protection key may be read or written.

If CR4.PKE = 0, or if 4-level paging is not active, the processor does not associate linear addresses with protection keys and does not use the access-control mechanism described in this section. In either of these cases, a reference in Section 4.6.1 to a user-mode address with a protection key should be considered a reference to any user-mode address.

The PKRU register (protection key rights for user pages) is a 32-bit register with the following format: for each  $i$  ( $0 \leq i \leq 15$ ), PKRU[2*i*] is the **access-disable bit** for protection key  $i$  (AD*i*); PKRU[2*i*+1] is the **write-disable bit** for protection key  $i$  (WD*i*).

Software can use the RDPKRU and WRPKRU instructions with ECX = 0 to read and write PKRU. In addition, the PKRU register is XSAVE-managed state and can thus be read and written by instructions in the XSAVE feature set. See Chapter 13, "Managing State Using the XSAVE Feature Set," of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1* for more information about the XSAVE feature set.

How a linear address's protection key controls access to the address depends on the mode of the linear address:

- A linear address's protection key controls only data accesses to the address. It does not in any way affect instructions fetched from the address.
- The protection key of a supervisor-mode address is ignored and does not control data accesses to the address. Because of this, Section 4.6.1 does not refer to protection keys when specifying the access rights for supervisor-mode addresses.
- Use of the protection key  $i$  of a user-mode address depends on the value of the PKRU register:
  - If AD*i* = 1, no data accesses are permitted.
  - If WD*i* = 1, permission may be denied to certain data write accesses:
    - User-mode write accesses are not permitted.
    - Supervisor-mode write accesses are not permitted if CR0.WP = 1. (If CR0.WP = 0, WD*i* does not affect supervisor-mode write accesses to user-mode addresses with protection key  $i$ .)

## 4.7 PAGE-FAULT EXCEPTIONS

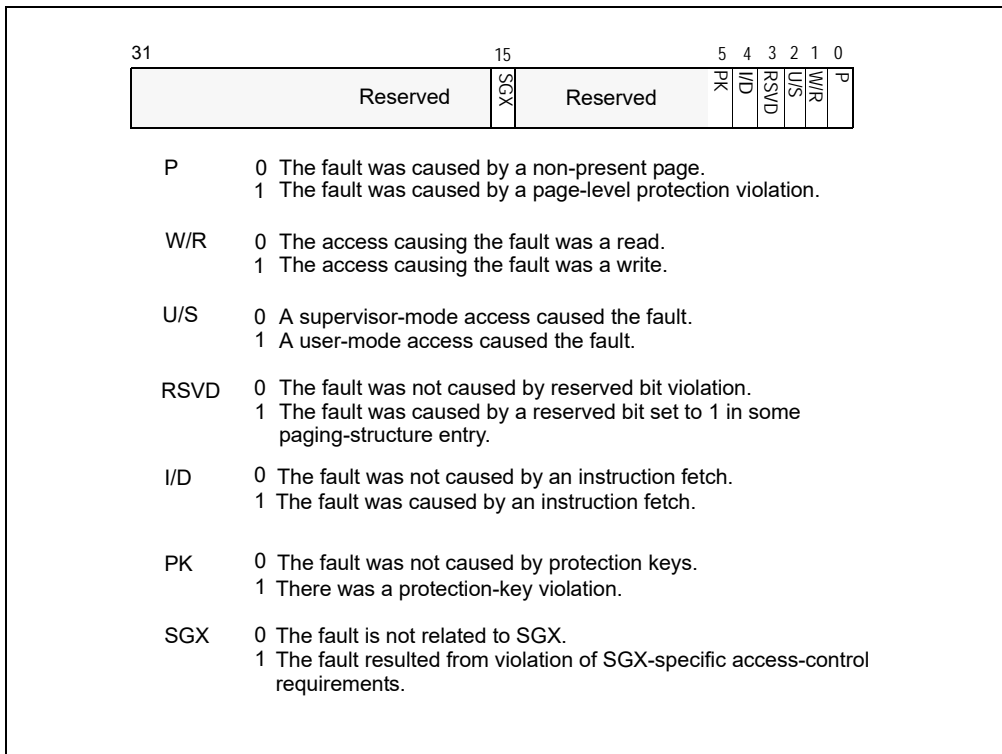
Accesses using linear addresses may cause **page-fault exceptions** (#PF; exception 14). An access to a linear address may cause a page-fault exception for either of two reasons: (1) there is no translation for the linear address; or (2) there is a translation for the linear address, but its access rights do not permit the access.

## PAGING

As noted in Section 4.3, Section 4.4.2, and Section 4.5, there is no translation for a linear address if the translation process for that address would use a paging-structure entry in which the P flag (bit 0) is 0 or one that sets a reserved bit. If there is a translation for a linear address, its access rights are determined as specified in Section 4.6.

When Intel® Software Guard Extensions (Intel® SGX) are enabled, the processor may deliver exception 14 for reasons unrelated to paging. See Section 37.3, “Access-control Requirements” and Section 37.19, “Enclave Page Cache Map (EPCM)” in Chapter 37, “Enclave Access Control and Data Structures.” Such an exception is called an **SGX-induced page fault**. The processor uses the error code to distinguish SGX-induced page faults from ordinary page faults.

Figure 4-12 illustrates the error code that the processor provides on delivery of a page-fault exception. The following items explain how the bits in the error code describe the nature of the page-fault exception:



**Figure 4-12. Page-Fault Error Code**

- **P flag (bit 0).**  
This flag is 0 if there is no translation for the linear address because the P flag was 0 in one of the paging-structure entries used to translate that address.
- **W/R (bit 1).**  
If the access causing the page-fault exception was a write, this flag is 1; otherwise, it is 0. This flag describes the access causing the page-fault exception, not the access rights specified by paging.
- **U/S (bit 2).**  
If a user-mode access caused the page-fault exception, this flag is 1; it is 0 if a supervisor-mode access did so. This flag describes the access causing the page-fault exception, not the access rights specified by paging. User-mode and supervisor-mode accesses are defined in Section 4.6.



- **RSVD flag (bit 3).**  
This flag is 1 if there is no translation for the linear address because a reserved bit was set in one of the paging-structure entries used to translate that address. (Because reserved bits are not checked in a paging-structure entry whose P flag is 0, bit 3 of the error code can be set only if bit 0 is also set.<sup>1</sup>)  
Bits reserved in the paging-structure entries are reserved for future functionality. Software developers should be aware that such bits may be used in the future and that a paging-structure entry that causes a page-fault exception on one processor might not do so in the future.
- **I/D flag (bit 4).**  
This flag is 1 if (1) the access causing the page-fault exception was an instruction fetch; and (2) either (a) CR4.SMEP = 1; or (b) both (i) CR4.PAE = 1 (either PAE paging or 4-level paging is in use); and (ii) IA32\_EFER.NXE = 1. Otherwise, the flag is 0. This flag describes the access causing the page-fault exception, not the access rights specified by paging.
- **PK flag (bit 5).**  
This flag is 1 if (1) IA32\_EFER.LMA = CR4.PKE = 1; (2) the access causing the page-fault exception was a data access; (3) the linear address was a user-mode address with protection key *i*; and (5) the PKRU register (see Section 4.6.2) is such that either (a) AD<sub>*i*</sub> = 1; or (b) the following all hold: (i) WD<sub>*i*</sub> = 1; (ii) the access is a write access; and (iii) either CR0.WP = 1 or the access causing the page-fault exception was a user-mode access.
- **SGX flag (bit 15).**  
This flag is 1 if the exception is unrelated to paging and resulted from violation of SGX-specific access-control requirements. Because such a violation can occur only if there is no ordinary page fault, this flag is set only if the P flag (bit 0) is 1 and the RSVD flag (bit 3) and the PK flag (bit 5) are both 0.

Page-fault exceptions occur only due to an attempt to use a linear address. Failures to load the PDPTTE registers with PAE paging (see Section 4.4.1) cause general-protection exceptions (#GP(0)) and not page-fault exceptions.

## 4.8 ACCESSED AND DIRTY FLAGS

For any paging-structure entry that is used during linear-address translation, bit 5 is the **accessed flag**.<sup>2</sup> For paging-structure entries that map a page (as opposed to referencing another paging structure), bit 6 is the **dirty flag**. These flags are provided for use by memory-management software to manage the transfer of pages and paging structures into and out of physical memory.

Whenever the processor uses a paging-structure entry as part of linear-address translation, it sets the accessed flag in that entry (if it is not already set).

Whenever there is a write to a linear address, the processor sets the dirty flag (if it is not already set) in the paging-structure entry that identifies the final physical address for the linear address (either a PTE or a paging-structure entry in which the PS flag is 1).

Memory-management software may clear these flags when a page or a paging structure is initially loaded into physical memory. These flags are “sticky,” meaning that, once set, the processor does not clear them; only software can clear them.

A processor may cache information from the paging-structure entries in TLBs and paging-structure caches (see Section 4.10). This fact implies that, if software changes an accessed flag or a dirty flag from 1 to 0, the processor might not set the corresponding bit in memory on a subsequent access using an affected linear address (see Section 4.10.4.3). See Section 4.10.4.2 for how software can ensure that these bits are updated as desired.

### NOTE

The accesses used by the processor to set these flags may or may not be exposed to the processor’s self-modifying code detection logic. If the processor is executing code from the same

- 
1. Some past processors had errata for some page faults that occur when there is no translation for the linear address because the P flag was 0 in one of the paging-structure entries used to translate that address. Due to these errata, some such page faults produced error codes that cleared bit 0 (P flag) and set bit 3 (RSVD flag).
  2. With PAE paging, the PDPTTEs are not used during linear-address translation but only to load the PDPTTE registers for some executions of the MOV CR instruction (see Section 4.4.1). For this reason, the PDPTTEs do not contain accessed flags with PAE paging.

memory area that is being used for the paging structures, the setting of these flags may or may not result in an immediate change to the executing code stream.

## 4.9 PAGING AND MEMORY TYPING

The **memory type** of a memory access refers to the type of caching used for that access. Chapter 11, “Memory Cache Control” provides many details regarding memory typing in the Intel-64 and IA-32 architectures. This section describes how paging contributes to the determination of memory typing.

The way in which paging contributes to memory typing depends on whether the processor supports the **Page Attribute Table (PAT)**; see Section 11.12).<sup>1</sup> Section 4.9.1 and Section 4.9.2 explain how paging contributes to memory typing depending on whether the PAT is supported.

### 4.9.1 Paging and Memory Typing When the PAT is Not Supported (Pentium Pro and Pentium II Processors)

#### NOTE

The PAT is supported on all processors that support 4-level paging. Thus, this section applies only to 32-bit paging and PAE paging.

If the PAT is not supported, paging contributes to memory typing in conjunction with the memory-type range registers (MTRRs) as specified in Table 11-6 in Section 11.5.2.1.

For any access to a physical address, the table combines the memory type specified for that physical address by the MTRRs with a PCD value and a PWT value. The latter two values are determined as follows:

- For an access to a PDE with 32-bit paging, the PCD and PWT values come from CR3.
- For an access to a PDE with PAE paging, the PCD and PWT values come from the relevant PDPTTE register.
- For an access to a PTE, the PCD and PWT values come from the relevant PDE.
- For an access to the physical address that is the translation of a linear address, the PCD and PWT values come from the relevant PTE (if the translation uses a 4-KByte page) or the relevant PDE (otherwise).
- With PAE paging, the UC memory type is used when loading the PDPTTEs (see Section 4.4.1).

### 4.9.2 Paging and Memory Typing When the PAT is Supported (Pentium III and More Recent Processor Families)

If the PAT is supported, paging contributes to memory typing in conjunction with the PAT and the memory-type range registers (MTRRs) as specified in Table 11-7 in Section 11.5.2.2.

The PAT is a 64-bit MSR (IA32\_PAT; MSR index 277H) comprising eight (8) 8-bit entries (entry *i* comprises bits  $8i+7:8i$  of the MSR).

---

1. The PAT is supported on Pentium III and more recent processor families. See Section 4.1.4 for how to determine whether the PAT is supported.

For any access to a physical address, the table combines the memory type specified for that physical address by the MTRRs with a memory type selected from the PAT. Table 11-11 in Section 11.12.3 specifies how a memory type is selected from the PAT. Specifically, it comes from entry  $i$  of the PAT, where  $i$  is defined as follows:

- For an access to an entry in a paging structure whose address is in CR3 (e.g., the PML4 table with 4-level paging):
  - For 4-level paging with  $CR4.PCIDE = 1$ ,  $i = 0$ .
  - Otherwise,  $i = 2*PCD + PWT$ , where the PCD and PWT values come from CR3.
- For an access to a PDE with PAE paging,  $i = 2*PCD + PWT$ , where the PCD and PWT values come from the relevant PDPTTE register.
- For an access to a paging-structure entry X whose address is in another paging-structure entry Y,  $i = 2*PCD + PWT$ , where the PCD and PWT values come from Y.
- For an access to the physical address that is the translation of a linear address,  $i = 4*PAT + 2*PCD + PWT$ , where the PAT, PCD, and PWT values come from the relevant PTE (if the translation uses a 4-KByte page), the relevant PDE (if the translation uses a 2-MByte page or a 4-MByte page), or the relevant PDPTTE (if the translation uses a 1-GByte page).
- With PAE paging, the WB memory type is used when loading the PDPTTEs (see Section 4.4.1).<sup>1</sup>

### 4.9.3 Caching Paging-Related Information about Memory Typing

A processor may cache information from the paging-structure entries in TLBs and paging-structure caches (see Section 4.10). These structures may include information about memory typing. The processor may use memory-typing information from the TLBs and paging-structure caches instead of from the paging structures in memory.

This fact implies that, if software modifies a paging-structure entry to change the memory-typing bits, the processor might not use that change for a subsequent translation using that entry or for access to an affected linear address. See Section 4.10.4.2 for how software can ensure that the processor uses the modified memory typing.

## 4.10 CACHING TRANSLATION INFORMATION

The Intel-64 and IA-32 architectures may accelerate the address-translation process by caching data from the paging structures on the processor. Because the processor does not ensure that the data that it caches are always consistent with the structures in memory, it is important for software developers to understand how and when the processor may cache such data. They should also understand what actions software can take to remove cached data that may be inconsistent and when it should do so. This section provides software developers information about the relevant processor operation.

Section 4.10.1 introduces process-context identifiers (PCIDs), which a logical processor may use to distinguish information cached for different linear-address spaces. Section 4.10.2 and Section 4.10.3 describe how the processor may cache information in translation lookaside buffers (TLBs) and paging-structure caches, respectively. Section 4.10.4 explains how software can remove inconsistent cached information by invalidating portions of the TLBs and paging-structure caches. Section 4.10.5 describes special considerations for multiprocessor systems.

### 4.10.1 Process-Context Identifiers (PCIDs)

Process-context identifiers (PCIDs) are a facility by which a logical processor may cache information for multiple linear-address spaces. The processor may retain cached information when software switches to a different linear-address space with a different PCID (e.g., by loading CR3; see Section 4.10.4.1 for details).

---

1. Some older IA-32 processors used the UC memory type when loading the PDPTTEs. Some processors may use the UC memory type if  $CR0.CD = 1$  or if the MTRRs are disabled. These behaviors are model-specific and not architectural.

A PCID is a 12-bit identifier. Non-zero PCIDs are enabled by setting the PCIDE flag (bit 17) of CR4. If CR4.PCIDE = 0, the current PCID is always 000H; otherwise, the current PCID is the value of bits 11:0 of CR3. Not all processors allow CR4.PCIDE to be set to 1; see Section 4.1.4 for how to determine whether this is allowed.

The processor ensures that CR4.PCIDE can be 1 only in IA-32e mode (thus, 32-bit paging and PAE paging use only PCID 000H). In addition, software can change CR4.PCIDE from 0 to 1 only if CR3[11:0] = 000H. These requirements are enforced by the following limitations on the MOV CR instruction:

- MOV to CR4 causes a general-protection exception (#GP) if it would change CR4.PCIDE from 0 to 1 and either IA32\_EFER.LMA = 0 or CR3[11:0] ≠ 000H.
- MOV to CR0 causes a general-protection exception if it would clear CR0.PG to 0 while CR4.PCIDE = 1.

When a logical processor creates entries in the TLBs (Section 4.10.2) and paging-structure caches (Section 4.10.3), it associates those entries with the current PCID. When using entries in the TLBs and paging-structure caches to translate a linear address, a logical processor uses only those entries associated with the current PCID (see Section 4.10.2.4 for an exception).

If CR4.PCIDE = 0, a logical processor does not cache information for any PCID other than 000H. This is because (1) if CR4.PCIDE = 0, the logical processor will associate any newly cached information with the current PCID, 000H; and (2) if MOV to CR4 clears CR4.PCIDE, all cached information is invalidated (see Section 4.10.4.1).

## NOTE

In revisions of this manual that were produced when no processors allowed CR4.PCIDE to be set to 1, Section 4.10 discussed the caching of translation information without any reference to PCIDs. While the section now refers to PCIDs in its specification of this caching, this documentation change is not intended to imply any change to the behavior of processors that do not allow CR4.PCIDE to be set to 1.

## 4.10.2 Translation Lookaside Buffers (TLBs)

A processor may cache information about the translation of linear addresses in translation lookaside buffers (TLBs). In general, TLBs contain entries that map page numbers to page frames; these terms are defined in Section 4.10.2.1. Section 4.10.2.2 describes how information may be cached in TLBs, and Section 4.10.2.3 gives details of TLB usage. Section 4.10.2.4 explains the global-page feature, which allows software to indicate that certain translations should receive special treatment when cached in the TLBs.

### 4.10.2.1 Page Numbers, Page Frames, and Page Offsets

Section 4.3, Section 4.4.2, and Section 4.5 give details of how the different paging modes translate linear addresses to physical addresses. Specifically, the upper bits of a linear address (called the **page number**) determine the upper bits of the physical address (called the **page frame**); the lower bits of the linear address (called the **page offset**) determine the lower bits of the physical address. The boundary between the page number and the page offset is determined by the **page size**. Specifically:

- 32-bit paging:
  - If the translation does not use a PTE (because CR4.PSE = 1 and the PS flag is 1 in the PDE used), the page size is 4 MBytes and the page number comprises bits 31:22 of the linear address.
  - If the translation does use a PTE, the page size is 4 KBytes and the page number comprises bits 31:12 of the linear address.
- PAE paging:
  - If the translation does not use a PTE (because the PS flag is 1 in the PDE used), the page size is 2 MBytes and the page number comprises bits 31:21 of the linear address.
  - If the translation does use a PTE, the page size is 4 KBytes and the page number comprises bits 31:12 of the linear address.

- 4-level paging:
  - If the translation does not use a PDE (because the PS flag is 1 in the PDPTE used), the page size is 1 GByte and the page number comprises bits 47:30 of the linear address.
  - If the translation does use a PDE but does not use a PTE (because the PS flag is 1 in the PDE used), the page size is 2 MBytes and the page number comprises bits 47:21 of the linear address.
  - If the translation does use a PTE, the page size is 4 KBytes and the page number comprises bits 47:12 of the linear address.

#### 4.10.2.2 Caching Translations in TLBs

The processor may accelerate the paging process by caching individual translations in **translation lookaside buffers (TLBs)**. Each entry in a TLB is an individual translation. Each translation is referenced by a page number. It contains the following information from the paging-structure entries used to translate linear addresses with the page number:

- The physical address corresponding to the page number (the page frame).
- The access rights from the paging-structure entries used to translate linear addresses with the page number (see Section 4.6):
  - The logical-AND of the R/W flags.
  - The logical-AND of the U/S flags.
  - The logical-OR of the XD flags (necessary only if IA32\_EFER.NXE = 1).
  - The protection key (necessary only with 4-level paging and CR4.PKE = 1).
- Attributes from a paging-structure entry that identifies the final page frame for the page number (either a PTE or a paging-structure entry in which the PS flag is 1):
  - The dirty flag (see Section 4.8).
  - The memory type (see Section 4.9).

(TLB entries may contain other information as well. A processor may implement multiple TLBs, and some of these may be for special purposes, e.g., only for instruction fetches. Such special-purpose TLBs may not contain some of this information if it is not necessary. For example, a TLB used only for instruction fetches need not contain information about the R/W and dirty flags.)

As noted in Section 4.10.1, any TLB entries created by a logical processor are associated with the current PCID.

Processors need not implement any TLBs. Processors that do implement TLBs may invalidate any TLB entry at any time. Software should not rely on the existence of TLBs or on the retention of TLB entries.

#### 4.10.2.3 Details of TLB Use

Because the TLBs cache entries only for linear addresses with translations, there can be a TLB entry for a page number only if the P flag is 1 and the reserved bits are 0 in each of the paging-structure entries used to translate that page number. In addition, the processor does not cache a translation for a page number unless the accessed flag is 1 in each of the paging-structure entries used during translation; before caching a translation, the processor sets any of these accessed flags that is not already 1.

The processor may cache translations required for prefetches and for accesses that are a result of speculative execution that would never actually occur in the executed code path.

If the page number of a linear address corresponds to a TLB entry associated with the current PCID, the processor may use that TLB entry to determine the page frame, access rights, and other attributes for accesses to that linear address. In this case, the processor may not actually consult the paging structures in memory. The processor may retain a TLB entry unmodified even if software subsequently modifies the relevant paging-structure entries in memory. See Section 4.10.4.2 for how software can ensure that the processor uses the modified paging-structure entries.

If the paging structures specify a translation using a page larger than 4 KBytes, some processors may cache multiple smaller-page TLB entries for that translation. Each such TLB entry would be associated with a page

number corresponding to the smaller page size (e.g., bits 47:12 of a linear address with 4-level paging), even though part of that page number (e.g., bits 20:12) is part of the offset with respect to the page specified by the paging structures. The upper bits of the physical address in such a TLB entry are derived from the physical address in the PDE used to create the translation, while the lower bits come from the linear address of the access for which the translation is created. There is no way for software to be aware that multiple translations for smaller pages have been used for a large page. For example, an execution of INVLPG for a linear address on such a page invalidates any and all smaller-page TLB entries for the translation of any linear address on that page.

If software modifies the paging structures so that the page size used for a 4-KByte range of linear addresses changes, the TLBs may subsequently contain multiple translations for the address range (one for each page size). A reference to a linear address in the address range may use any of these translations. Which translation is used may vary from one execution to another, and the choice may be implementation-specific.

#### 4.10.2.4 Global Pages

The Intel-64 and IA-32 architectures also allow for **global pages** when the PGE flag (bit 7) is 1 in CR4. If the G flag (bit 8) is 1 in a paging-structure entry that maps a page (either a PTE or a paging-structure entry in which the PS flag is 1), any TLB entry cached for a linear address using that paging-structure entry is considered to be **global**. Because the G flag is used only in paging-structure entries that map a page, and because information from such entries is not cached in the paging-structure caches, the global-page feature does not affect the behavior of the paging-structure caches.

A logical processor may use a global TLB entry to translate a linear address, even if the TLB entry is associated with a PCID different from the current PCID.

### 4.10.3 Paging-Structure Caches

In addition to the TLBs, a processor may cache other information about the paging structures in memory.

#### 4.10.3.1 Caches for Paging Structures

A processor may support any or all of the following paging-structure caches:

- **PML4 cache** (4-level paging only). Each PML4-cache entry is referenced by a 9-bit value and is used for linear addresses for which bits 47:39 have that value. The entry contains information from the PML4E used to translate such linear addresses:
  - The physical address from the PML4E (the address of the page-directory-pointer table).
  - The value of the R/W flag of the PML4E.
  - The value of the U/S flag of the PML4E.
  - The value of the XD flag of the PML4E.
  - The values of the PCD and PWT flags of the PML4E.

The following items detail how a processor may use the PML4 cache:

- If the processor has a PML4-cache entry for a linear address, it may use that entry when translating the linear address (instead of the PML4E in memory).
- The processor does not create a PML4-cache entry unless the P flag is 1 and all reserved bits are 0 in the PML4E in memory.
- The processor does not create a PML4-cache entry unless the accessed flag is 1 in the PML4E in memory; before caching a translation, the processor sets the accessed flag if it is not already 1.
- The processor may create a PML4-cache entry even if there are no translations for any linear address that might use that entry (e.g., because the P flags are 0 in all entries in the referenced page-directory-pointer table).
- If the processor creates a PML4-cache entry, the processor may retain it unmodified even if software subsequently modifies the corresponding PML4E in memory.

- **PDPTE cache (4-level paging only).**<sup>1</sup> Each PDPTE-cache entry is referenced by an 18-bit value and is used for linear addresses for which bits 47:30 have that value. The entry contains information from the PML4E and PDPTE used to translate such linear addresses:
  - The physical address from the PDPTTE (the address of the page directory). (No PDPTTE-cache entry is created for a PDPTTE that maps a 1-GByte page.)
  - The logical-AND of the R/W flags in the PML4E and the PDPTTE.
  - The logical-AND of the U/S flags in the PML4E and the PDPTTE.
  - The logical-OR of the XD flags in the PML4E and the PDPTTE.
  - The values of the PCD and PWT flags of the PDPTTE.

The following items detail how a processor may use the PDPTTE cache:

- If the processor has a PDPTTE-cache entry for a linear address, it may use that entry when translating the linear address (instead of the PML4E and the PDPTTE in memory).
  - The processor does not create a PDPTTE-cache entry unless the P flag is 1, the PS flag is 0, and the reserved bits are 0 in the PML4E and the PDPTTE in memory.
  - The processor does not create a PDPTTE-cache entry unless the accessed flags are 1 in the PML4E and the PDPTTE in memory; before caching a translation, the processor sets any accessed flags that are not already 1.
  - The processor may create a PDPTTE-cache entry even if there are no translations for any linear address that might use that entry.
  - If the processor creates a PDPTTE-cache entry, the processor may retain it unmodified even if software subsequently modifies the corresponding PML4E or PDPTTE in memory.
- **PDE cache.** The use of the PDE cache depends on the paging mode:
    - For 32-bit paging, each PDE-cache entry is referenced by a 10-bit value and is used for linear addresses for which bits 31:22 have that value.
    - For PAE paging, each PDE-cache entry is referenced by an 11-bit value and is used for linear addresses for which bits 31:21 have that value.
    - For 4-level paging, each PDE-cache entry is referenced by a 27-bit value and is used for linear addresses for which bits 47:21 have that value.

A PDE-cache entry contains information from the PML4E, PDPTTE, and PDE used to translate the relevant linear addresses (for 32-bit paging and PAE paging, only the PDE applies):

- The physical address from the PDE (the address of the page table). (No PDE-cache entry is created for a PDE that maps a page.)
- The logical-AND of the R/W flags in the PML4E, PDPTTE, and PDE.
- The logical-AND of the U/S flags in the PML4E, PDPTTE, and PDE.
- The logical-OR of the XD flags in the PML4E, PDPTTE, and PDE.
- The values of the PCD and PWT flags of the PDE.

---

1. With PAE paging, the PDPTTEs are stored in internal, non-architectural registers. The operation of these registers is described in Section 4.4.1 and differs from that described here.

The following items detail how a processor may use the PDE cache (references below to PML4Es and PDPTEs apply only to 4-level paging):

- If the processor has a PDE-cache entry for a linear address, it may use that entry when translating the linear address (instead of the PML4E, the PDPTE, and the PDE in memory).
- The processor does not create a PDE-cache entry unless the P flag is 1, the PS flag is 0, and the reserved bits are 0 in the PML4E, the PDPTE, and the PDE in memory.
- The processor does not create a PDE-cache entry unless the accessed flag is 1 in the PML4E, the PDPTE, and the PDE in memory; before caching a translation, the processor sets any accessed flags that are not already 1.
- The processor may create a PDE-cache entry even if there are no translations for any linear address that might use that entry.
- If the processor creates a PDE-cache entry, the processor may retain it unmodified even if software subsequently modifies the corresponding PML4E, the PDPTE, or the PDE in memory.

Information from a paging-structure entry can be included in entries in the paging-structure caches for other paging-structure entries referenced by the original entry. For example, if the R/W flag is 0 in a PML4E, then the R/W flag will be 0 in any PDPTE-cache entry for a PDPTE from the page-directory-pointer table referenced by that PML4E. This is because the R/W flag of each such PDPTE-cache entry is the logical-AND of the R/W flags in the appropriate PML4E and PDPTE.

The paging-structure caches contain information only from paging-structure entries that reference other paging structures (and not those that map pages). Because the G flag is not used in such paging-structure entries, the global-page feature does not affect the behavior of the paging-structure caches.

The processor may create entries in paging-structure caches for translations required for prefetches and for accesses that are a result of speculative execution that would never actually occur in the executed code path.

As noted in Section 4.10.1, any entries created in paging-structure caches by a logical processor are associated with the current PCID.

A processor may or may not implement any of the paging-structure caches. Software should rely on neither their presence nor their absence. The processor may invalidate entries in these caches at any time. Because the processor may create the cache entries at the time of translation and not update them following subsequent modifications to the paging structures in memory, software should take care to invalidate the cache entries appropriately when causing such modifications. The invalidation of TLBs and the paging-structure caches is described in Section 4.10.4.

### 4.10.3.2 Using the Paging-Structure Caches to Translate Linear Addresses

When a linear address is accessed, the processor uses a procedure such as the following to determine the physical address to which it translates and whether the access should be allowed:

- If the processor finds a TLB entry that is for the page number of the linear address and that is associated with the current PCID (or which is global), it may use the physical address, access rights, and other attributes from that entry.
- If the processor does not find a relevant TLB entry, it may use the upper bits of the linear address to select an entry from the PDE cache that is associated with the current PCID (Section 4.10.3.1 indicates which bits are used in each paging mode). It can then use that entry to complete the translation process (locating a PTE, etc.) as if it had traversed the PDE (and, for 4-level paging, the PDPTE and PML4E) corresponding to the PDE-cache entry.
- The following items apply when 4-level paging is used:
  - If the processor does not find a relevant TLB entry or a relevant PDE-cache entry, it may use bits 47:30 of the linear address to select an entry from the PDPTE cache that is associated with the current PCID. It can then use that entry to complete the translation process (locating a PDE, etc.) as if it had traversed the PDPTE and the PML4E corresponding to the PDPTE-cache entry.
  - If the processor does not find a relevant TLB entry, a relevant PDE-cache entry, or a relevant PDPTE-cache entry, it may use bits 47:39 of the linear address to select an entry from the PML4 cache that is associated



with the current PCID. It can then use that entry to complete the translation process (locating a PDPTE, etc.) as if it had traversed the corresponding PML4E.

(Any of the above steps would be skipped if the processor does not support the cache in question.)

If the processor does not find a TLB or paging-structure-cache entry for the linear address, it uses the linear address to traverse the entire paging-structure hierarchy, as described in Section 4.3, Section 4.4.2, and Section 4.5.

### 4.10.3.3 Multiple Cached Entries for a Single Paging-Structure Entry

The paging-structure caches and TLBs may contain multiple entries associated with a single PCID and with information derived from a single paging-structure entry. The following items give some examples for 4-level paging:

- Suppose that two PML4Es contain the same physical address and thus reference the same page-directory-pointer table. Any PDPTE in that table may result in two PDPTE-cache entries, each associated with a different set of linear addresses. Specifically, suppose that the  $n_1^{\text{th}}$  and  $n_2^{\text{th}}$  entries in the PML4 table contain the same physical address. This implies that the physical address in the  $m^{\text{th}}$  PDPTE in the page-directory-pointer table would appear in the PDPTE-cache entries associated with both  $p_1$  and  $p_2$ , where  $(p_1 \gg 9) = n_1$ ,  $(p_2 \gg 9) = n_2$ , and  $(p_1 \& 1\text{FFH}) = (p_2 \& 1\text{FFH}) = m$ . This is because both PDPTE-cache entries use the same PDPTE, one resulting from a reference from the  $n_1^{\text{th}}$  PML4E and one from the  $n_2^{\text{th}}$  PML4E.
- Suppose that the first PML4E (i.e., the one in position 0) contains the physical address X in CR3 (the physical address of the PML4 table). This implies the following:
  - Any PML4-cache entry associated with linear addresses with 0 in bits 47:39 contains address X.
  - Any PDPTE-cache entry associated with linear addresses with 0 in bits 47:30 contains address X. This is because the translation for a linear address for which the value of bits 47:30 is 0 uses the value of bits 47:39 (0) to locate a page-directory-pointer table at address X (the address of the PML4 table). It then uses the value of bits 38:30 (also 0) to find address X again and to store that address in the PDPTE-cache entry.
  - Any PDE-cache entry associated with linear addresses with 0 in bits 47:21 contains address X for similar reasons.
  - Any TLB entry for page number 0 (associated with linear addresses with 0 in bits 47:12) translates to page frame  $X \gg 12$  for similar reasons.

The same PML4E contributes its address X to all these cache entries because the self-referencing nature of the entry causes it to be used as a PML4E, a PDPTE, a PDE, and a PTE.

## 4.10.4 Invalidation of TLBs and Paging-Structure Caches

As noted in Section 4.10.2 and Section 4.10.3, the processor may create entries in the TLBs and the paging-structure caches when linear addresses are translated, and it may retain these entries even after the paging structures used to create them have been modified. To ensure that linear-address translation uses the modified paging structures, software should take action to invalidate any cached entries that may contain information that has since been modified.

### 4.10.4.1 Operations that Invalidate TLBs and Paging-Structure Caches

The following instructions invalidate entries in the TLBs and the paging-structure caches:

- INVLPG. This instruction takes a single operand, which is a linear address. The instruction invalidates any TLB entries that are for a page number corresponding to the linear address and that are associated with the current PCID. It also invalidates any global TLB entries with that page number, regardless of PCID (see Section 4.10.2.4).<sup>1</sup> INVLPG also invalidates all entries in all paging-structure caches associated with the current PCID, regardless of the linear addresses to which they correspond.

---

1. If the paging structures map the linear address using a page larger than 4 KBytes and there are multiple TLB entries for that page (see Section 4.10.2.3), the instruction invalidates all of them.

- INVPCID. The operation of this instruction is based on instruction operands, called the INVPCID type and the INVPCID descriptor. Four INVPCID types are currently defined:
  - Individual-address. If the INVPCID type is 0, the logical processor invalidates mappings—except global translations—associated with the PCID specified in the INVPCID descriptor and that would be used to translate the linear address specified in the INVPCID descriptor.<sup>1</sup> (The instruction may also invalidate global translations, as well as mappings associated with other PCIDs and for other linear addresses.)
  - Single-context. If the INVPCID type is 1, the logical processor invalidates all mappings—except global translations—associated with the PCID specified in the INVPCID descriptor. (The instruction may also invalidate global translations, as well as mappings associated with other PCIDs.)
  - All-context, including globals. If the INVPCID type is 2, the logical processor invalidates mappings—including global translations—associated with all PCIDs.
  - All-context. If the INVPCID type is 3, the logical processor invalidates mappings—except global translations—associated with all PCIDs. (The instruction may also invalidate global translations.)

See Chapter 3 of the *Intel 64 and IA-32 Architecture Software Developer's Manual, Volume 2A* for details of the INVPCID instruction.

- MOV to CR0. The instruction invalidates all TLB entries (including global entries) and all entries in all paging-structure caches (for all PCIDs) if it changes the value of CR0.PG from 1 to 0.
- MOV to CR3. The behavior of the instruction depends on the value of CR4.PCIDE:
  - If CR4.PCIDE = 0, the instruction invalidates all TLB entries associated with PCID 000H except those for global pages. It also invalidates all entries in all paging-structure caches associated with PCID 000H.
  - If CR4.PCIDE = 1 and bit 63 of the instruction's source operand is 0, the instruction invalidates all TLB entries associated with the PCID specified in bits 11:0 of the instruction's source operand except those for global pages. It also invalidates all entries in all paging-structure caches associated with that PCID. It is not required to invalidate entries in the TLBs and paging-structure caches that are associated with other PCIDs.
  - If CR4.PCIDE = 1 and bit 63 of the instruction's source operand is 1, the instruction is not required to invalidate any TLB entries or entries in paging-structure caches.
- MOV to CR4. The behavior of the instruction depends on the bits being modified:
  - The instruction invalidates all TLB entries (including global entries) and all entries in all paging-structure caches (for all PCIDs) if (1) it changes the value of CR4.PGE;<sup>2</sup> or (2) it changes the value of the CR4.PCIDE from 1 to 0.
  - The instruction invalidates all TLB entries and all entries in all paging-structure caches for the current PCID if (1) it changes the value of CR4.PAE; or (2) it changes the value of CR4.SMEP from 0 to 1.
- Task switch. If a task switch changes the value of CR3, it invalidates all TLB entries associated with PCID 000H except those for global pages. It also invalidates all entries in all paging-structure caches associated with PCID 000H.<sup>3</sup>
- VMX transitions. See Section 4.11.1.

The processor is always free to invalidate additional entries in the TLBs and paging-structure caches. The following are some examples:

- INVLPG may invalidate TLB entries for pages other than the one corresponding to its linear-address operand. It may invalidate TLB entries and paging-structure-cache entries associated with PCIDs other than the current PCID.

---

1. If the paging structures map the linear address using a page larger than 4 KBytes and there are multiple TLB entries for that page (see Section 4.10.2.3), the instruction invalidates all of them.

2. If CR4.PGE is changing from 0 to 1, there were no global TLB entries before the execution; if CR4.PGE is changing from 1 to 0, there will be no global TLB entries after the execution.

3. Task switches do not occur in IA-32e mode and thus cannot occur with 4-level paging. Since CR4.PCIDE can be set only with 4-level paging, task switches occur only with CR4.PCIDE = 0.

- INVPID may invalidate TLB entries for pages other than the one corresponding to the specified linear address. It may invalidate TLB entries and paging-structure-cache entries associated with PCIDs other than the specified PCID.
- MOV to CR0 may invalidate TLB entries even if CR0.PG is not changing. For example, this may occur if either CR0.CD or CR0.NW is modified.
- MOV to CR3 may invalidate TLB entries for global pages. If CR4.PCIDE = 1 and bit 63 of the instruction's source operand is 0, it may invalidate TLB entries and entries in the paging-structure caches associated with PCIDs other than the PCID it is establishing. It may invalidate entries if CR4.PCIDE = 1 and bit 63 of the instruction's source operand is 1.
- MOV to CR4 may invalidate TLB entries when changing CR4.PSE or when changing CR4.SMEP from 1 to 0.
- On a processor supporting Hyper-Threading Technology, invalidations performed on one logical processor may invalidate entries in the TLBs and paging-structure caches used by other logical processors.

(Other instructions and operations may invalidate entries in the TLBs and the paging-structure caches, but the instructions identified above are recommended.)

In addition to the instructions identified above, page faults invalidate entries in the TLBs and paging-structure caches. In particular, a page-fault exception resulting from an attempt to use a linear address will invalidate any TLB entries that are for a page number corresponding to that linear address and that are associated with the current PCID. It also invalidates all entries in the paging-structure caches that would be used for that linear address and that are associated with the current PCID.<sup>1</sup> These invalidations ensure that the page-fault exception will not recur (if the faulting instruction is re-executed) if it would not be caused by the contents of the paging structures in memory (and if, therefore, it resulted from cached entries that were not invalidated after the paging structures were modified in memory).

As noted in Section 4.10.2, some processors may choose to cache multiple smaller-page TLB entries for a translation specified by the paging structures to use a page larger than 4 KBytes. There is no way for software to be aware that multiple translations for smaller pages have been used for a large page. The INVLPG instruction and page faults provide the same assurances that they provide when a single TLB entry is used: they invalidate all TLB entries corresponding to the translation specified by the paging structures.

#### 4.10.4.2 Recommended Invalidation

The following items provide some recommendations regarding when software should perform invalidations:

- If software modifies a paging-structure entry that maps a page (rather than referencing another paging structure), it should execute INVLPG for any linear address with a page number whose translation uses that paging-structure entry.<sup>2</sup>  
(If the paging-structure entry may be used in the translation of different page numbers — see Section 4.10.3.3 — software should execute INVLPG for linear addresses with each of those page numbers; alternatively, it could use MOV to CR3 or MOV to CR4.)
- If software modifies a paging-structure entry that references another paging structure, it may use one of the following approaches depending upon the types and number of translations controlled by the modified entry:
  - Execute INVLPG for linear addresses with each of the page numbers with translations that would use the entry. However, if no page numbers that would use the entry have translations (e.g., because the P flags are 0 in all entries in the paging structure referenced by the modified entry), it remains necessary to execute INVLPG at least once.
  - Execute MOV to CR3 if the modified entry controls no global pages.
  - Execute MOV to CR4 to modify CR4.PGE.
- If CR4.PCIDE = 1 and software modifies a paging-structure entry that does not map a page or in which the G flag (bit 8) is 0, additional steps are required if the entry may be used for PCIDs other than the current one. Any one of the following suffices:

---

1. Unlike INVLPG, page faults need not invalidate **all** entries in the paging-structure caches, only those that would be used to translate the faulting linear address.

2. One execution of INVLPG is sufficient even for a page with size greater than 4 KBytes.

- Execute MOV to CR4 to modify CR4.PGE, either immediately or before again using any of the affected PCIDs. For example, software could use different (previously unused) PCIDs for the processes that used the affected PCIDs.
- For each affected PCID, execute MOV to CR3 to make that PCID current (and to load the address of the appropriate PML4 table). If the modified entry controls no global pages and bit 63 of the source operand to MOV to CR3 was 0, no further steps are required. Otherwise, execute INVLPG for linear addresses with each of the page numbers with translations that would use the entry; if no page numbers that would use the entry have translations, execute INVLPG at least once.
- If software using PAE paging modifies a PDPTE, it should reload CR3 with the register's current value to ensure that the modified PDPTE is loaded into the corresponding PDPTE register (see Section 4.4.1).
- If the nature of the paging structures is such that a single entry may be used for multiple purposes (see Section 4.10.3.3), software should perform invalidations for all of these purposes. For example, if a single entry might serve as both a PDE and PTE, it may be necessary to execute INVLPG with two (or more) linear addresses, one that uses the entry as a PDE and one that uses it as a PTE. (Alternatively, software could use MOV to CR3 or MOV to CR4.)
- As noted in Section 4.10.2, the TLBs may subsequently contain multiple translations for the address range if software modifies the paging structures so that the page size used for a 4-KByte range of linear addresses changes. A reference to a linear address in the address range may use any of these translations.  
Software wishing to prevent this uncertainty should not write to a paging-structure entry in a way that would change, for any linear address, both the page size and either the page frame, access rights, or other attributes. It can instead use the following algorithm: first clear the P flag in the relevant paging-structure entry (e.g., PDE); then invalidate any translations for the affected linear addresses (see above); and then modify the relevant paging-structure entry to set the P flag and establish modified translation(s) for the new page size.
- Software should clear bit 63 of the source operand to a MOV to CR3 instruction that establishes a PCID that had been used earlier for a different linear-address space (e.g., with a different value in bits 51:12 of CR3). This ensures invalidation of any information that may have been cached for the previous linear-address space.  
This assumes that both linear-address spaces use the same global pages and that it is thus not necessary to invalidate any global TLB entries. If that is not the case, software should invalidate those entries by executing MOV to CR4 to modify CR4.PGE.

#### 4.10.4.3 Optional Invalidation

The following items describe cases in which software may choose not to invalidate and the potential consequences of that choice:

- If a paging-structure entry is modified to change the P flag from 0 to 1, no invalidation is necessary. This is because no TLB entry or paging-structure cache entry is created with information from a paging-structure entry in which the P flag is 0.<sup>1</sup>
- If a paging-structure entry is modified to change the accessed flag from 0 to 1, no invalidation is necessary (assuming that an invalidation was performed the last time the accessed flag was changed from 1 to 0). This is because no TLB entry or paging-structure cache entry is created with information from a paging-structure entry in which the accessed flag is 0.
- If a paging-structure entry is modified to change the R/W flag from 0 to 1, failure to perform an invalidation may result in a "spurious" page-fault exception (e.g., in response to an attempted write access) but no other adverse behavior. Such an exception will occur at most once for each affected linear address (see Section 4.10.4.1).
- If CR4.SMEP = 0 and a paging-structure entry is modified to change the U/S flag from 0 to 1, failure to perform an invalidation may result in a "spurious" page-fault exception (e.g., in response to an attempted user-mode access) but no other adverse behavior. Such an exception will occur at most once for each affected linear address (see Section 4.10.4.1).
- If a paging-structure entry is modified to change the XD flag from 1 to 0, failure to perform an invalidation may result in a "spurious" page-fault exception (e.g., in response to an attempted instruction fetch) but no other

---

1. If it is also the case that no invalidation was performed the last time the P flag was changed from 1 to 0, the processor may use a TLB entry or paging-structure cache entry that was created when the P flag had earlier been 1.

adverse behavior. Such an exception will occur at most once for each affected linear address (see Section 4.10.4.1).

- If a paging-structure entry is modified to change the accessed flag from 1 to 0, failure to perform an invalidation may result in the processor not setting that bit in response to a subsequent access to a linear address whose translation uses the entry. Software cannot interpret the bit being clear as an indication that such an access has not occurred.
- If software modifies a paging-structure entry that identifies the final physical address for a linear address (either a PTE or a paging-structure entry in which the PS flag is 1) to change the dirty flag from 1 to 0, failure to perform an invalidation may result in the processor not setting that bit in response to a subsequent write to a linear address whose translation uses the entry. Software cannot interpret the bit being clear as an indication that such a write has not occurred.
- The read of a paging-structure entry in translating an address being used to fetch an instruction may appear to execute before an earlier write to that paging-structure entry if there is no serializing instruction between the write and the instruction fetch. Note that the invalidating instructions identified in Section 4.10.4.1 are all serializing instructions.
- Section 4.10.3.3 describes situations in which a single paging-structure entry may contain information cached in multiple entries in the paging-structure caches. Because all entries in these caches are invalidated by any execution of INVLPG, it is not necessary to follow the modification of such a paging-structure entry by executing INVLPG multiple times solely for the purpose of invalidating these multiple cached entries. (It may be necessary to do so to invalidate multiple TLB entries.)

#### 4.10.4.4 Delayed Invalidation

Required invalidations may be delayed under some circumstances. Software developers should understand that, between the modification of a paging-structure entry and execution of the invalidation instruction recommended in Section 4.10.4.2, the processor may use translations based on either the old value or the new value of the paging-structure entry. The following items describe some of the potential consequences of delayed invalidation:

- If a paging-structure entry is modified to change the P flag from 1 to 0, an access to a linear address whose translation is controlled by this entry may or may not cause a page-fault exception.
- If a paging-structure entry is modified to change the R/W flag from 0 to 1, write accesses to linear addresses whose translation is controlled by this entry may or may not cause a page-fault exception.
- If a paging-structure entry is modified to change the U/S flag from 0 to 1, user-mode accesses to linear addresses whose translation is controlled by this entry may or may not cause a page-fault exception.
- If a paging-structure entry is modified to change the XD flag from 1 to 0, instruction fetches from linear addresses whose translation is controlled by this entry may or may not cause a page-fault exception.

As noted in Section 8.1.1, an x87 instruction or an SSE instruction that accesses data larger than a quadword may be implemented using multiple memory accesses. If such an instruction stores to memory and invalidation has been delayed, some of the accesses may complete (writing to memory) while another causes a page-fault exception.<sup>1</sup> In this case, the effects of the completed accesses may be visible to software even though the overall instruction caused a fault.

In some cases, the consequences of delayed invalidation may not affect software adversely. For example, when freeing a portion of the linear-address space (by marking paging-structure entries “not present”), invalidation using INVLPG may be delayed if software does not re-allocate that portion of the linear-address space or the memory that had been associated with it. However, because of speculative execution (or errant software), there may be accesses to the freed portion of the linear-address space before the invalidations occur. In this case, the following can happen:

- Reads can occur to the freed portion of the linear-address space. Therefore, invalidation should not be delayed for an address range that has read side effects.
- The processor may retain entries in the TLBs and paging-structure caches for an extended period of time. Software should not assume that the processor will not use entries associated with a linear address simply because time has passed.

---

1. If the accesses are to different pages, this may occur even if invalidation has not been delayed.

- As noted in Section 4.10.3.1, the processor may create an entry in a paging-structure cache even if there are no translations for any linear address that might use that entry. Thus, if software has marked “not present” all entries in a page table, the processor may subsequently create a PDE-cache entry for the PDE that references that page table (assuming that the PDE itself is marked “present”).
- If software attempts to write to the freed portion of the linear-address space, the processor might not generate a page fault. (Such an attempt would likely be the result of a software error.) For that reason, the page frames previously associated with the freed portion of the linear-address space should not be reallocated for another purpose until the appropriate invalidations have been performed.

#### 4.10.5 Propagation of Paging-Structure Changes to Multiple Processors

As noted in Section 4.10.4, software that modifies a paging-structure entry may need to invalidate entries in the TLBs and paging-structure caches that were derived from the modified entry before it was modified. In a system containing more than one logical processor, software must account for the fact that there may be entries in the TLBs and paging-structure caches of logical processors other than the one used to modify the paging-structure entry. The process of propagating the changes to a paging-structure entry is commonly referred to as “TLB shutdown.”

TLB shutdown can be done using memory-based semaphores and/or interprocessor interrupts (IPI). The following items describe a simple but inefficient example of a TLB shutdown algorithm for processors supporting the Intel-64 and IA-32 architectures:

1. Begin barrier: Stop all but one logical processor; that is, cause all but one to execute the HLT instruction or to enter a spin loop.
2. Allow the active logical processor to change the necessary paging-structure entries.
3. Allow all logical processors to perform invalidations appropriate to the modifications to the paging-structure entries.
4. Allow all logical processors to resume normal operation.

Alternative, performance-optimized, TLB shutdown algorithms may be developed; however, software developers must take care to ensure that the following conditions are met:

- All logical processors that are using the paging structures that are being modified must participate and perform appropriate invalidations after the modifications are made.
- If the modifications to the paging-structure entries are made before the barrier or if there is no barrier, the operating system must ensure one of the following: (1) that the affected linear-address range is not used between the time of modification and the time of invalidation; or (2) that it is prepared to deal with the consequences of the affected linear-address range being used during that period. For example, if the operating system does not allow pages being freed to be reallocated for another purpose until after the required invalidations, writes to those pages by errant software will not unexpectedly modify memory that is in use.
- Software must be prepared to deal with reads, instruction fetches, and prefetch requests to the affected linear-address range that are a result of speculative execution that would never actually occur in the executed code path.

When multiple logical processors are using the same linear-address space at the same time, they must coordinate before any request to modify the paging-structure entries that control that linear-address space. In these cases, the barrier in the TLB shutdown routine may not be required. For example, when freeing a range of linear addresses, some other mechanism can assure no logical processor is using that range before the request to free it is made. In this case, a logical processor freeing the range can clear the P flags in the PTEs associated with the range, free the physical page frames associated with the range, and then signal the other logical processors using that linear-address space to perform the necessary invalidations. All the affected logical processors must complete their invalidations before the linear-address range and the physical page frames previously associated with that range can be reallocated.

## 4.11 INTERACTIONS WITH VIRTUAL-MACHINE EXTENSIONS (VMX)

The architecture for virtual-machine extensions (VMX) includes features that interact with paging. Section 4.11.1 discusses ways in which VMX-specific control transfers, called VMX transitions specially affect paging. Section 4.11.2 gives an overview of VMX features specifically designed to support address translation.

### 4.11.1 VMX Transitions

The VMX architecture defines two control transfers called **VM entries** and **VM exits**; collectively, these are called **VMX transitions**. VM entries and VM exits are described in detail in Chapter 26 and Chapter 27, respectively, in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C*. The following items identify paging-related details:

- VMX transitions modify the CR0 and CR4 registers and the IA32\_EFER MSR concurrently. For this reason, they allow transitions between paging modes that would not otherwise be possible:
  - VM entries allow transitions from 4-level paging directly to either 32-bit paging or PAE paging.
  - VM exits allow transitions from either 32-bit paging or PAE paging directly to 4-level paging.
- VMX transitions that result in PAE paging load the PDPTTE registers (see Section 4.4.1) as follows:
  - VM entries load the PDPTTE registers either from the physical address being loaded into CR3 or from the virtual-machine control structure (VMCS); see Section 26.3.2.4.
  - VM exits load the PDPTTE registers from the physical address being loaded into CR3; see Section 27.5.4.
- VMX transitions invalidate the TLBs and paging-structure caches based on certain control settings. See Section 26.3.2.5 and Section 27.5.5 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C*.

### 4.11.2 VMX Support for Address Translation

Chapter 28, “VMX Support for Address Translation,” in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C* describe two features of the virtual-machine extensions (VMX) that interact directly with paging. These are **virtual-processor identifiers (VPIDs)** and the **extended page table mechanism (EPT)**.

VPIDs provide a way for software to identify to the processor the address spaces for different “virtual processors.” The processor may use this identification to maintain concurrently information for multiple address spaces in its TLBs and paging-structure caches, even when non-zero PCIDs are not being used. See Section 28.1 for details.

When EPT is in use, the addresses in the paging-structures are not used as physical addresses to access memory and memory-mapped I/O. Instead, they are treated as **guest-physical** addresses and are translated through a set of EPT paging structures to produce physical addresses. EPT can also specify its own access rights and memory typing; these are used on conjunction with those specified in this chapter. See Section 28.2 for more information.

Both VPIDs and EPT may change the way that a processor maintains information in TLBs and paging structure caches and the ways in which software can manage that information. Some of the behaviors documented in Section 4.10 may change. See Section 28.3 for details.

## 4.12 USING PAGING FOR VIRTUAL MEMORY

With paging, portions of the linear-address space need not be mapped to the physical-address space; data for the unmapped addresses can be stored externally (e.g., on disk). This method of mapping the linear-address space is referred to as virtual memory or demand-paged virtual memory.

Paging divides the linear address space into fixed-size pages that can be mapped into the physical-address space and/or external storage. When a program (or task) references a linear address, the processor uses paging to translate the linear address into a corresponding physical address if such an address is defined.

If the page containing the linear address is not currently mapped into the physical-address space, the processor generates a page-fault exception as described in Section 4.7. The handler for page-fault exceptions typically

directs the operating system or executive to load data for the unmapped page from external storage into physical memory (perhaps writing a different page from physical memory out to external storage in the process) and to map it using paging (by updating the paging structures). When the page has been loaded into physical memory, a return from the exception handler causes the instruction that generated the exception to be restarted.

Paging differs from segmentation through its use of fixed-size pages. Unlike segments, which usually are the same size as the code or data structures they hold, pages have a fixed size. If segmentation is the only form of address translation used, a data structure present in physical memory will have all of its parts in memory. If paging is used, a data structure can be partly in memory and partly in disk storage.

### 4.13 MAPPING SEGMENTS TO PAGES

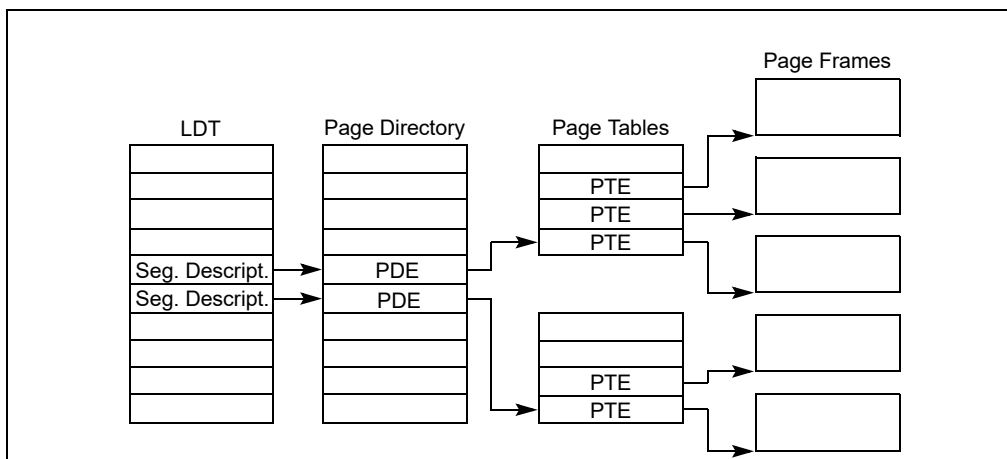
The segmentation and paging mechanisms provide support for a wide variety of approaches to memory management. When segmentation and paging are combined, segments can be mapped to pages in several ways. To implement a flat (unsegmented) addressing environment, for example, all the code, data, and stack modules can be mapped to one or more large segments (up to 4-GBytes) that share same range of linear addresses (see Figure 3-2 in Section 3.2.2). Here, segments are essentially invisible to applications and the operating-system or executive. If paging is used, the paging mechanism can map a single linear-address space (contained in a single segment) into virtual memory. Alternatively, each program (or task) can have its own large linear-address space (contained in its own segment), which is mapped into virtual memory through its own paging structures.

Segments can be smaller than the size of a page. If one of these segments is placed in a page which is not shared with another segment, the extra memory is wasted. For example, a small data structure, such as a 1-Byte semaphore, occupies 4 KBytes if it is placed in a page by itself. If many semaphores are used, it is more efficient to pack them into a single page.

The Intel-64 and IA-32 architectures do not enforce correspondence between the boundaries of pages and segments. A page can contain the end of one segment and the beginning of another. Similarly, a segment can contain the end of one page and the beginning of another.

Memory-management software may be simpler and more efficient if it enforces some alignment between page and segment boundaries. For example, if a segment which can fit in one page is placed in two pages, there may be twice as much paging overhead to support access to that segment.

One approach to combining paging and segmentation that simplifies memory-management software is to give each segment its own page table, as shown in Figure 4-13. This convention gives the segment a single entry in the page directory, and this entry provides the access control information for paging the entire segment.



**Figure 4-13. Memory Management Convention That Assigns a Page Table to Each Segment**



## 10. Updates to Chapter 14, Volume 3B

Change bars show changes to Chapter 14 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2*.

-----

Changes to this chapter: Addition of new power management features and minor typo corrections.

This chapter describes facilities of Intel 64 and IA-32 architecture used for power management and thermal monitoring.

### 14.1 ENHANCED INTEL SPEEDSTEP® TECHNOLOGY

Enhanced Intel SpeedStep® Technology was introduced in the Pentium M processor. The technology enables the management of processor power consumption via performance state transitions. These states are defined as discrete operating points associated with different voltages and frequencies.

Enhanced Intel SpeedStep Technology differs from previous generations of Intel SpeedStep® Technology in two ways:

- Centralization of the control mechanism and software interface in the processor by using model-specific registers.
- Reduced hardware overhead; this permits more frequent performance state transitions.

Previous generations of the Intel SpeedStep Technology require processors to be a deep sleep state, holding off bus master transfers for the duration of a performance state transition. Performance state transitions under the Enhanced Intel SpeedStep Technology are discrete transitions to a new target frequency.

Support is indicated by CPUID, using ECX feature bit 07. Enhanced Intel SpeedStep Technology is enabled by setting IA32\_MISC\_ENABLE MSR, bit 16. On reset, bit 16 of IA32\_MISC\_ENABLE MSR is cleared.

#### 14.1.1 Software Interface For Initiating Performance State Transitions

State transitions are initiated by writing a 16-bit value to the IA32\_PERF\_CTL register, see Figure 14-2. If a transition is already in progress, transition to a new value will subsequently take effect.

Reads of IA32\_PERF\_CTL determine the last targeted operating point. The current operating point can be read from IA32\_PERF\_STATUS. IA32\_PERF\_STATUS is updated dynamically.

The 16-bit encoding that defines valid operating points is model-specific. Applications and performance tools are not expected to use either IA32\_PERF\_CTL or IA32\_PERF\_STATUS and should treat both as reserved. Performance monitoring tools can access model-specific events and report the occurrences of state transitions.

### 14.2 P-STATE HARDWARE COORDINATION


The Advanced Configuration and Power Interface (ACPI) defines performance states (P-states) that are used to facilitate system software's ability to manage processor power consumption. Different P-states correspond to different performance levels that are applied while the processor is actively executing instructions. Enhanced Intel SpeedStep Technology supports P-states by providing software interfaces that control the operating frequency and voltage of a processor.

With multiple processor cores residing in the same physical package, hardware dependencies may exist for a subset of logical processors on a platform. These dependencies may impose requirements that impact the coordination of P-state transitions. As a result, multi-core processors may require an OS to provide additional software support for coordinating P-state transitions for those subsets of logical processors.

ACPI firmware can choose to expose P-states as dependent and hardware-coordinated to OS power management (OSPM) policy. To support OSPMs, multi-core processors must have additional built-in support for P-state hardware coordination and feedback.

Intel 64 and IA-32 processors with dependent P-states amongst a subset of logical processors permit hardware coordination of P-states and provide a hardware-coordination feedback mechanism using IA32\_MPERF MSR and

IA32\_APERF MSR. See Figure 14-1 for an overview of the two 64-bit MSRs and the bullets below for a detailed description.



**Figure 14-1. IA32\_MPERF MSR and IA32\_APERF MSR for P-state Coordination**

- Use CPUID to check the P-State hardware coordination feedback capability bit. CPUID.06H.ECX[Bit 0] = 1 indicates IA32\_MPERF MSR and IA32\_APERF MSR are present.
- IA32\_MPERF MSR (E7H) increments in proportion to a fixed frequency, which is configured when the processor is booted.
- IA32\_APERF MSR (E8H) increments in proportion to actual performance, while accounting for hardware coordination of P-state and TM1/TM2; or software initiated throttling.
- The MSRs are per logical processor; they measure performance only when the targeted processor is in the C0 state.
- Only the IA32\_APERF/IA32\_MPERF ratio is architecturally defined; software should not attach meaning to the content of the individual of IA32\_APERF or IA32\_MPERF MSRs.
- When either MSR overflows, both MSRs are reset to zero and continue to increment.
- Both MSRs are full 64-bits counters. Each MSR can be written to independently. However, software should follow the guidelines illustrated in Example 14-1.

If P-states are exposed by the BIOS as hardware coordinated, software is expected to confirm processor support for P-state hardware coordination feedback and use the feedback mechanism to make P-state decisions. The OSPM is expected to either save away the current MSR values (for determination of the delta of the counter ratio at a later time) or reset both MSRs (execute WRMSR with 0 to these MSRs individually) at the start of the time window used for making the P-state decision. When not resetting the values, overflow of the MSRs can be detected by checking whether the new values read are less than the previously saved values.

Example 14-1 demonstrates steps for using the hardware feedback mechanism provided by IA32\_APERF MSR and IA32\_MPERF MSR to determine a target P-state.

#### Example 14-1. Determine Target P-state From Hardware Coordinated Feedback

```

DWORD PercentBusy; // Percentage of processor time not idle.
// Measure "PercentBusy" during previous sampling window.
// Typically, "PercentBusy" is measure over a time scale suitable for
// power management decisions
//
// RDMSR of MCNT and ACNT should be performed without delay.
// Software needs to exercise care to avoid delays between
// the two RDMSRs (for example, interrupts).
MCNT = RDMSR(IA32_MPERF);
ACNT = RDMSR(IA32_APERF);

// PercentPerformance indicates the percentage of the processor
// that is in use. The calculation is based on the PercentBusy,
// that is the percentage of processor time not idle and the P-state
// hardware coordinated feedback using the ACNT/MCNT ratio.
// Note that both values need to be calculated over the same

```

```

// time window.
    PercentPerformance = PercentBusy * (ACNT/MCNT);

// This example does not cover the additional logic or algorithms
// necessary to coordinate multiple logical processors to a target P-state.

TargetPstate = FindPstate(PercentPerformance);

if (TargetPstate ≠ currentPstate) {
    SetPState(TargetPstate);
}
// WRMSR of MCNT and ACNT should be performed without delay.
// Software needs to exercise care to avoid delays between
// the two WRMSRs (for example, interrupts).
WRMSR(IA32_MPERF, 0);
WRMSR(IA32_APERF, 0);

```

## 14.3 SYSTEM SOFTWARE CONSIDERATIONS AND OPPORTUNISTIC PROCESSOR PERFORMANCE OPERATION

An Intel 64 processor may support a form of processor operation that takes advantage of design headroom to opportunistically increase performance. The Intel<sup>®</sup> Turbo Boost Technology can convert thermal headroom into higher performance across multi-threaded and single-threaded workloads. The Intel<sup>®</sup> Dynamic Acceleration Technology feature can convert thermal headroom into higher performance if only one thread is active.

### 14.3.1 Intel<sup>®</sup> Dynamic Acceleration Technology

The Intel Core 2 Duo processor T 7700 introduces Intel Dynamic Acceleration Technology. Intel Dynamic Acceleration Technology takes advantage of thermal design headroom and opportunistically allows a single core to operate at a higher performance level when the operating system requests increased performance.

### 14.3.2 System Software Interfaces for Opportunistic Processor Performance Operation

Opportunistic processor performance operation, applicable to Intel Dynamic Acceleration Technology and Intel<sup>®</sup> Turbo Boost Technology, has the following characteristics:

- A transition from a normal state of operation (e.g. Intel Dynamic Acceleration Technology/Turbo mode disengaged) to a target state is not guaranteed, but may occur opportunistically after the corresponding enable mechanism is activated, the headroom is available and certain criteria are met.
- The opportunistic processor performance operation is generally transparent to most application software.
- System software (BIOS and Operating system) must be aware of hardware support for opportunistic processor performance operation and may need to temporarily disengage opportunistic processor performance operation when it requires more predictable processor operation.
- When opportunistic processor performance operation is engaged, the OS should use hardware coordination feedback mechanisms to prevent un-intended policy effects if it is activated during inappropriate situations.

#### 14.3.2.1 Discover Hardware Support and Enabling of Opportunistic Processor Performance Operation

If an Intel 64 processor has hardware support for opportunistic processor performance operation, the power-on default state of IA32\_MISC\_ENABLE[38] indicates the presence of such hardware support. For Intel 64 processors that support opportunistic processor performance operation, the default value is 1, indicating its presence. For processors that do not support opportunistic processor performance operation, the default value is 0. The power-

on default value of IA32\_MISC\_ENABLE[38] allows BIOS to detect the presence of hardware support of opportunistic processor performance operation.

IA32\_MISC\_ENABLE[38] is shared across all logical processors in a physical package. It is written by BIOS during platform initiation to enable/disable opportunistic processor performance operation in conjunction of OS power management capabilities, see Section 14.3.2.2. BIOS can set IA32\_MISC\_ENABLE[38] with 1 to disable opportunistic processor performance operation; it must clear the default value of IA32\_MISC\_ENABLE[38] to 0 to enable opportunistic processor performance operation. OS and applications must use CPUID leaf 06H if it needs to detect processors that have opportunistic processor performance operation enabled.

When CPUID is executed with EAX = 06H on input, Bit 1 of EAX in Leaf 06H (i.e. CPUID.06H:EAX[1]) indicates opportunistic processor performance operation, such as Intel Dynamic Acceleration Technology, has been enabled by BIOS.

Opportunistic processor performance operation can be disabled by setting bit 38 of IA32\_MISC\_ENABLE. This mechanism is intended for BIOS only. If IA32\_MISC\_ENABLE[38] is set, CPUID.06H:EAX[1] will return 0.

### 14.3.2.2 OS Control of Opportunistic Processor Performance Operation

There may be phases of software execution in which system software cannot tolerate the non-deterministic aspects of opportunistic processor performance operation. For example, when calibrating a real-time workload to make a CPU reservation request to the OS, it may be undesirable to allow the possibility of the processor delivering increased performance that cannot be sustained after the calibration phase.

System software can temporarily disengage opportunistic processor performance operation by setting bit 32 of the IA32\_PERF\_CTL MSR (0199H), using a read-modify-write sequence on the MSR. The opportunistic processor performance operation can be re-engaged by clearing bit 32 in IA32\_PERF\_CTL MSR, using a read-modify-write sequence. The DISENAGE bit in IA32\_PERF\_CTL is not reflected in bit 32 of the IA32\_PERF\_STATUS MSR (0198H), and it is not shared between logical processors in a physical package. In order for OS to engage Intel Dynamic Acceleration Technology/Turbo mode, the BIOS must:

- Enable opportunistic processor performance operation, as described in Section 14.3.2.1.
- Expose the operating points associated with Intel Dynamic Acceleration Technology/Turbo mode to the OS.




Figure 14-2. IA32\_PERF\_CTL Register

### 14.3.2.3 Required Changes to OS Power Management P-State Policy

Intel Dynamic Acceleration Technology and Intel Turbo Boost Technology can provide opportunistic performance greater than the performance level corresponding to the Processor Base frequency of the processor (see CPUID's processor frequency information). System software can use a pair of MSRs to observe performance feedback. Software must query for the presence of IA32\_APERF and IA32\_MPERF (see Section 14.2). The ratio between IA32\_APERF and IA32\_MPERF is architecturally defined and a value greater than unity indicates performance increase occurred during the observation period due to Intel Dynamic Acceleration Technology. Without incorporating such performance feedback, the target P-state evaluation algorithm can result in a non-optimal P-state target.

There are other scenarios under which OS power management may want to disable Intel Dynamic Acceleration Technology, some of these are listed below:

- When engaging ACPI defined passive thermal management, it may be more effective to disable Intel Dynamic Acceleration Technology for the duration of passive thermal management.
- When the user has indicated a policy preference of power savings over performance, OS power management may want to disable Intel Dynamic Acceleration Technology while that policy is in effect.

### 14.3.3 Intel® Turbo Boost Technology

Intel Turbo Boost Technology is supported in Intel Core i7 processors and Intel Xeon processors based on Intel® microarchitecture code name Nehalem. It uses the same principle of leveraging thermal headroom to dynamically increase processor performance for single-threaded and multi-threaded/multi-tasking environment. The programming interface described in Section 14.3.2 also applies to Intel Turbo Boost Technology.

### 14.3.4 Performance and Energy Bias Hint support

Intel 64 processors may support additional software hint to guide the hardware heuristic of power management features to favor increasing dynamic performance or conserve energy consumption.

Software can detect the processor's capability to support the performance-energy bias preference hint by examining bit 3 of ECX in CPUID leaf 6. The processor supports this capability if CPUID.06H:ECX.SETBH[bit 3] is set and it also implies the presence of a new architectural MSR called IA32\_ENERGY\_PERF\_BIAS (1B0H).

Software can program the lowest four bits of IA32\_ENERGY\_PERF\_BIAS MSR with a value from 0 - 15. The values represent a sliding scale, where a value of 0 (the default reset value) corresponds to a hint preference for highest performance and a value of 15 corresponds to the maximum energy savings. A value of 7 roughly translates into a hint to balance performance with energy consumption.




Figure 14-3. IA32\_ENERGY\_PERF\_BIAS Register

The layout of IA32\_ENERGY\_PERF\_BIAS is shown in Figure 14-3. The scope of IA32\_ENERGY\_PERF\_BIAS is per logical processor, which means that each of the logical processors in the package can be programmed with a different value. This may be especially important in virtualization scenarios, where the performance / energy requirements of one logical processor may differ from the other. Conflicting "hints" from various logical processors at higher hierarchy level will be resolved in favor of performance over energy savings.

Software can use whatever criteria it sees fit to program the MSR with an appropriate value. However, the value only serves as a hint to the hardware and the actual impact on performance and energy savings is model specific.

## 14.4 HARDWARE-CONTROLLED PERFORMANCE STATES (HWP)

Intel processors may contain support for Hardware-Controlled Performance States (HWP), which autonomously selects performance states while utilizing OS supplied performance guidance hints. The Enhanced Intel Speed-Step® Technology provides a means for the OS to control and monitor discrete frequency-based operating points via the IA32\_PERF\_CTL and IA32\_PERF\_STATUS MSRs.

In contrast, HWP is an implementation of the ACPI-defined Collaborative Processor Performance Control (CPPC), which specifies that the platform enumerates a continuous, abstract unit-less, performance value scale that is not tied to a specific performance state / frequency by definition. While the enumerated scale is roughly linear in terms of a delivered integer workload performance result, the OS is required to characterize the performance value range to comprehend the delivered performance for an applied workload.

When HWP is enabled, the processor autonomously selects performance states as deemed appropriate for the applied workload and with consideration of constraining hints that are programmed by the OS. These OS-provided hints include minimum and maximum performance limits, preference towards energy efficiency or performance, and the specification of a relevant workload history observation time window. The means for the OS to override HWP's autonomous selection of performance state with a specific desired performance target is also provided, however, the effective frequency delivered is subject to the result of energy efficiency and performance optimizations.

### 14.4.1 HWP Programming Interfaces

The programming interfaces provided by HWP include the following:

- The CPUID instruction allows software to discover the presence of HWP support in an Intel processor. Specifically, execute CPUID instruction with EAX=06H as input will return 5 bit flags covering the following aspects in bits 7 through 11 of CPUID.06H:EAX:
  - Availability of HWP baseline resource and capability, CPUID.06H:EAX[bit 7]: If this bit is set, HWP provides several new architectural MSRs: IA32\_PM\_ENABLE, IA32\_HWP\_CAPABILITIES, IA32\_HWP\_REQUEST, IA32\_HWP\_STATUS.
  - Availability of HWP Notification upon dynamic Guaranteed Performance change, CPUID.06H:EAX[bit 8]: If this bit is set, HWP provides IA32\_HWP\_INTERRUPT MSR to enable interrupt generation due to dynamic Performance changes and excursions.
  - Availability of HWP Activity window control, CPUID.06H:EAX[bit 9]: If this bit is set, HWP allows software to program activity window in the IA32\_HWP\_REQUEST MSR.
  - Availability of HWP energy/performance preference control, CPUID.06H:EAX[bit 10]: If this bit is set, HWP allows software to set an energy/performance preference hint in the IA32\_HWP\_REQUEST MSR.
  - Availability of HWP package level control, CPUID.06H:EAX[bit 11]: If this bit is set, HWP provides the IA32\_HWP\_REQUEST\_PKG MSR to convey OS Power Management's control hints for all logical processors in the physical package.

**Table 14-1. Architectural and Non-Architectural MSRs Related to HWP**

| Address | Architectural | Register Name                 | Description   |
|---------|---------------|-------------------------------|---|
| 770H    | Y             | IA32_PM_ENABLE                | Enable/Disable HWP.   |
| 771H    | Y             | IA32_HWP_CAPABILITIES         | Enumerates the HWP performance range (static and dynamic).  |
| 772H    | Y             | IA32_HWP_REQUEST_PKG          | Conveys OSPM's control hints (Min, Max, Activity Window, Energy Performance Preference, Desired) for all logical processor in the physical package. |
| 773H    | Y             | IA32_HWP_INTERRUPT            | Controls HWP native interrupt generation (Guaranteed Performance changes, excursions).  |
| 774H    | Y             | IA32_HWP_REQUEST              | Conveys OSPM's control hints (Min, Max, Activity Window, Energy Performance Preference, Desired) for a single logical processor.                    |
| 775H    | Y             | IA32_HWP_PECI_REQUEST_INFO    | Conveys embedded system controller requests to override some of the OS HWP Request settings via the Peci mechanism.                                 |
| 777H    | Y             | IA32_HWP_STATUS               | Status bits indicating changes to Guaranteed Performance and excursions to Minimum Performance.   |
| 19CH    | Y             | IA32_THERM_STATUS[bits 15:12] | Conveys reasons for performance excursions.   |
| 64EH    | N             | MSR_PPERF                     | Productive Performance Count.   |

- Additionally, HWP may provide a non-architectural MSR, MSR\_PPERF, which provides a quantitative metric to software of hardware's view of workload scalability. This hardware's view of workload scalability is implementation specific.

### 14.4.2 Enabling HWP

The layout of the IA32\_PM\_ENABLE MSR is shown in Figure 14-4. The bit fields are described below:




Figure 14-4. IA32\_PM\_ENABLE MSR

- **HWP\_ENABLE (bit 0, R/W1Once)** — Software sets this bit to enable HWP with autonomous selection of processor P-States. When set, the processor will disregard input from the legacy performance control interface (IA32\_PERF\_CTL). Note this bit can only be enabled once from the default value. Once set, writes to the HWP\_ENABLE bit are ignored. Only RESET will clear this bit. Default = zero (0).
- Bits 63:1 are reserved and must be zero.

After software queries CPUID and verifies the processor's support of HWP, system software can write 1 to IA32\_PM\_ENABLE.HWP\_ENABLE (bit 0) to enable hardware controlled performance states. The default value of IA32\_PM\_ENABLE MSR at power-on is 0, i.e. HWP is disabled.

Additional MSRs associated with HWP may only be accessed after HWP is enabled, with the exception of IA32\_HWP\_INTERRUPT and MSR\_PPERF. Accessing the IA32\_HWP\_INTERRUPT MSR requires only HWP is present as enumerated by CPUID but does not require enabling HWP.

IA32\_PM\_ENABLE is a package level MSR, i.e., writing to it from any logical processor within a package affects all logical processors within that package.

### 14.4.3 HWP Performance Range and Dynamic Capabilities

The OS reads the IA32\_HWP\_CAPABILITIES MSR to comprehend the limits of the HWP-managed performance range as well as the dynamic capability, which may change during processor operation. The enumerated performance range values reported by IA32\_HWP\_CAPABILITIES directly map to initial frequency targets (prior to workload-specific frequency optimizations of HWP). However the mapping is processor family specific.

The layout of the IA32\_HWP\_CAPABILITIES MSR is shown in Figure 14-5. The bit fields are described below:






Figure 14-5. IA32\_HWP\_CAPABILITIES Register

- **Highest\_Performance (bits 7:0, RO)** — Value for the maximum non-guaranteed performance level.
- **Guaranteed\_Performance (bits 15:8, RO)** — Current value for the guaranteed performance level. This value can change dynamically as a result of internal or external constraints, e.g. thermal or power limits.
- **Most\_Efficient\_Performance (bits 23:16, RO)** — Current value of the most efficient performance level. This value can change dynamically as a result of workload characteristics.
- **Lowest\_Performance (bits 31:24, RO)** — Value for the lowest performance level that software can program to IA32\_HWP\_REQUEST.
- Bits 63:32 are reserved and must be zero.

The value returned in the **Guaranteed\_Performance** field is hardware's best-effort approximation of the available performance given current operating constraints. Changes to the **Guaranteed\_Performance** value will primarily occur due to a shift in operational mode. This includes a power or other limit applied by an external agent, e.g. RAPL (see Figure 14.9.1), or the setting of a Configurable TDP level (see model-specific controls related to Programmable TDP Limit in Chapter 2, "Model-Specific Registers (MSRs)" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4.*). Notification of a change to the **Guaranteed\_Performance** occurs via interrupt (if configured) and the IA32\_HWP\_Status MSR. Changes to **Guaranteed\_Performance** are indicated when a macroscopically meaningful change in performance occurs i.e. sustained for greater than one second. Consequently, notification of a change in **Guaranteed\_Performance** will typically occur no more frequently than once per second. Rapid changes in platform configuration, e.g. docking / undocking, with corresponding changes to a Configurable TDP level could potentially cause more frequent notifications.

The value returned by the **Most\_Efficient\_Performance** field provides the OS with an indication of the practical lower limit for the IA32\_HWP\_REQUEST. The processor may not honor IA32\_HWP\_REQUEST.Maximum Performance settings below this value.

## 14.4.4 Managing HWP

### 14.4.4.1 IA32\_HWP\_REQUEST MSR (Address: 0x774 Logical Processor Scope)

Typically, the operating system controls HWP operation for each logical processor via the writing of control hints / constraints to the IA32\_HWP\_REQUEST MSR. The layout of the IA32\_HWP\_REQUEST MSR is shown in Figure 14-6. The bit fields are described below Figure 14-6.

Operating systems can control HWP by writing both IA32\_HWP\_REQUEST and IA32\_HWP\_REQUEST\_PKG MSRs (see Section 14.4.4.2). Five valid bits within the IA32\_HWP\_REQUEST MSR let the operating system flexibly select which of its five hint / constraint fields should be derived by the processor from the IA32\_HWP\_REQUEST MSR and which should be derived from the IA32\_HWP\_REQUEST\_PKG MSR. These five valid bits are supported if CPUID[6].EAX[17] is set.

When the IA32\_HWP\_REQUEST MSR Package Control bit is set, any valid bit that is NOT set indicates to the processor to use the respective field value from the IA32\_HWP\_REQUEST\_PKG MSR. Otherwise, the values are derived from the IA32\_HWP\_REQUEST MSR. The valid bits are ignored when the IA32\_HWP\_REQUEST MSR Package Control bit is zero.




Figure 14-6. IA32\_HWP\_REQUEST Register

- **Minimum\_Performance (bits 7:0, RW)** — Conveys a hint to the HWP hardware. The OS programs the minimum performance hint to achieve the required quality of service (QoS) or to meet a service level agreement (SLA) as needed. Note that an excursion below the level specified is possible due to hardware constraints. The default value of this field is IA32\_HWP\_CAPABILITIES.Lowest\_Performance.
- **Maximum\_Performance (bits 15:8, RW)** — Conveys a hint to the HWP hardware. The OS programs this field to limit the maximum performance that is expected to be supplied by the HWP hardware. Excursions above the limit requested by OS are possible due to hardware coordination between the processor cores and other components in the package. The default value of this field is IA32\_HWP\_CAPABILITIES.Highest\_Performance.
- **Desired\_Performance (bits 23:16, RW)** — Conveys a hint to the HWP hardware. When set to zero, hardware autonomous selection determines the performance target. When set to a non-zero value (between the range of Lowest\_Performance and Highest\_Performance of IA32\_HWP\_CAPABILITIES) conveys an explicit performance request hint to the hardware; effectively disabling HW Autonomous selection. The Desired\_Performance input is non-constraining in terms of Performance and Energy Efficiency optimizations, which are independently controlled. The default value of this field is 0.
- **Energy\_Performance\_Preference (bits 31:24, RW)** — Conveys a hint to the HWP hardware. The OS may write a range of values from 0 (performance preference) to 0FFH (energy efficiency preference) to influence the rate of performance increase /decrease and the result of the hardware's energy efficiency and performance optimizations. The default value of this field is 80H. Note: If CPUID.06H:EAX[bit 10] indicates that this field is not supported, HWP uses the value of the IA32\_ENERGY\_PERF\_BIAS MSR to determine the energy efficiency / performance preference.
- **Activity\_Window (bits 41:32, RW)** — Conveys a hint to the HWP hardware specifying a moving workload history observation window for performance/frequency optimizations. If 0, the hardware will determine the appropriate window size. When writing a non-zero value to this field, this field is encoded in the format of bits 38:32 as a 7-bit mantissa and bits 41:39 as a 3-bit exponent value in powers of 10. The resultant value is in microseconds. Thus, the minimal/maximum activity window size is 1 microsecond/1270 seconds. Combined with the Energy\_Performance\_Preference input, Activity\_Window influences the rate of performance increase

/ decrease. This non-zero hint only has meaning when `Desired_Performance = 0`. The default value of this field is 0.

- **Package\_Control (bit 42, RW)** — When set, causes this logical processor's `IA32_HWP_REQUEST` control inputs to be derived from the `IA32_HWP_REQUEST_PKG` MSR.
- Bits 58:43 are reserved and must be zero.
- **Activity\_Window\_Valid (bit 59, RW)** — When set, indicates to the processor to derive the Activity Window field value from the `IA32_HWP_REQUEST` MSR even if the package control bit is set. Otherwise, derive it from the `IA32_HWP_REQUEST_PKG` MSR. The default value of this field is 0.
- **EPP\_Valid (bit 60, RW)** — When set, indicates to the processor to derive the EPP field value from the `IA32_HWP_REQUEST` MSR even if the package control bit is set. Otherwise, derive it from the `IA32_HWP_REQUEST_PKG` MSR. The default value of this field is 0.
- **Desired\_Valid (bit 61, RW)** — When set, indicates to the processor to derive the Desired Performance field value from the `IA32_HWP_REQUEST` MSR even if the package control bit is set. Otherwise, derive it from the `IA32_HWP_REQUEST_PKG` MSR. The default value of this field is 0.
- **Maximum\_Valid (bit 62, RW)** — When set, indicates to the processor to derive the Maximum Performance field value from the `IA32_HWP_REQUEST` MSR even if the package control bit is set. Otherwise, derive it from the `IA32_HWP_REQUEST_PKG` MSR. The default value of this field is 0.
- **Minimum\_Valid (bit 63, RW)** — When set, indicates to the processor to derive the Minimum Performance field value from the `IA32_HWP_REQUEST` MSR even if the package control bit is set. Otherwise, derive it from the `IA32_HWP_REQUEST_PKG` MSR. The default value of this field is 0.

The HWP hardware clips and resolves the field values as necessary to the valid range. Reads return the last value written not the clipped values.

Processors may support a subset of `IA32_HWP_REQUEST` fields as indicated by `CPUID`. Reads of non-supported fields will return 0. Writes to non-supported fields are ignored.

The OS may override HWP's autonomous selection of performance state with a specific performance target by setting the `Desired_Performance` field to a non-zero value, however, the effective frequency delivered is subject to the result of energy efficiency and performance optimizations, which are influenced by the Energy Performance Preference field.

Software may disable all hardware optimizations by setting `Minimum_Performance = Maximum_Performance` (subject to package coordination).

Note: The processor may run below the `Minimum_Performance` level due to hardware constraints including: power, thermal, and package coordination constraints. The processor may also run below the `Minimum_Performance` level for short durations (few milliseconds) following C-state exit, and when Hardware Duty Cycling (see Section 14.5) is enabled.

When the `IA32_HWP_REQUEST` MSR is set to fast access mode, writes of this MSR are posted, i.e., the `WRMSR` instruction retires before the data reaches its destination within the processor. It may retire even before all preceding IA stores are globally visible, i.e., it is not an architecturally serializing instruction anymore (no store fence). A new `CPUID` bit indicates this new characteristic of the `IA32_HWP_REQUEST` MSR (see Section 14.4.8 for additional details).

#### 14.4.4.2 IA32\_HWP\_REQUEST\_PKG MSR (Address: 0x772 Package Scope)




Figure 14-7. IA32\_HWP\_REQUEST\_PKG Register

The structure of the IA32\_HWP\_REQUEST\_PKG MSR (package-level) is identical to the IA32\_HWP\_REQUEST MSR with the exception of the the Package Control bit field and the five valid bit fields, which do not exist in the IA32\_HWP\_REQUEST\_PKG MSR. Field values written to this MSR apply to all logical processors within the physical package with the exception of logical processors whose IA32\_HWP\_REQUEST.Package Control field is clear (zero). Single P-state Control mode is only supported when IA32\_HWP\_REQUEST\_PKG is not supported.

#### 14.4.4.3 IA32\_HWP\_PECI\_REQUEST\_INFO MSR (Address 0x775 Package Scope)

When an embedded system controller is integrated in the platform, it can override some of the OS HWP Request settings via the PECI mechanism. PECI initiated settings take precedence over the relevant fields in the IA32\_HWP\_REQUEST MSR and in the IA32\_HWP\_REQUEST\_PKG MSR, irrespective of the Package Control bit or the Valid Bit values described above. PECI can independently control each of: Minimum Performance, Maximum Performance and EPP fields. This MSR contains both the PECI induced values and the control bits that indicate whether the embedded controller actually set the processor to use the respective value.

PECI override is supported if CPUID[6].EAX[16] is set.




Figure 14-8. IA32\_HWP\_PECI\_REQUEST\_INFO MSR

The layout of the IA32\_HWP\_PECI\_REQUEST\_INFO MSR is shown in Figure 14-8. This MSR is writable by the embedded controller but is read-only by software executing on the CPU. This MSR has Package scope. The bit fields are described below:

- **Minimum\_Performance (bits 7:0, RO)** — Used by the OS to read the latest value of Peci minimum performance input.
- **Maximum\_Performance (bits 15:8, RO)** — Used by the OS to read the latest value of Peci maximum performance input.
- Bits 23:16 are reserved and must be zero.
- **Energy\_Performance\_Preference (bits 31:24, RO)** — Used by the OS to read the latest value of Peci energy performance preference input.
- Bits 59:32 are reserved and must be zero.
- **EPP\_Peci\_Override (bit 60, RO)** — Indicates whether Peci is currently overriding the Energy Performance Preference input. If set(1), Peci is overriding the Energy Performance Preference input. If clear(0), OS has control over Energy Performance Preference input.
- Bit 61 is reserved and must be zero.
- **Max\_Peci\_Override (bit 62, RO)** — Indicates whether Peci is currently overriding the Maximum Performance input. If set(1), Peci is overriding the Maximum Performance input. If clear(0), OS has control over Maximum Performance input.
- **Min\_Peci\_Override (bit 63, RO)** — Indicates whether Peci is currently overriding the Minimum Performance input. If set(1), Peci is overriding the Minimum Performance input. If clear(0), OS has control over Minimum Performance input.

### HWP Request Field Hierarchical Resolution

HWP Request field resolution is fed by three MSRs: IA32\_HWP\_REQUEST, IA32\_HWP\_REQUEST\_PKG and IA32\_HWP\_PECI\_REQUEST\_INFO. The flow that the processor goes through to resolve which field value is chosen is shown below.

For each of the two HWP Request fields; Desired and Activity Window:

```
If IA32_HWP_REQUEST.PACKAGE_CONTROL = 1 and IA32_HWP_REQUEST.<field> valid bit = 0
    Resolved Field Value = IA32_HWP_REQUEST_PKG.<field>
Else
    Resolved Field Value = IA32_HWP_REQUEST.<field>
```

For each of the three HWP Request fields; Min, Max and EPP:

```
If IA32_HWP_PECI_REQUEST_INFO.<field> Peci Override bit = 1
    Resolved Field Value = IA32_HWP_PECI_REQUEST_INFO.<field>
Else if IA32_HWP_REQUEST.PACKAGE_CONTROL = 1 and IA32_HWP_REQUEST.<field> valid bit = 0
    Resolved Field Value = IA32_HWP_REQUEST_PKG.<field>
Else
    Resolved Field Value = IA32_HWP_REQUEST.<field>
```

### 14.4.5 HWP Feedback

The processor provides several types of feedback to the OS during HWP operation.

The IA32\_MPERF MSR and IA32\_APERF MSR mechanism (see Section 14.2) allows the OS to calculate the resultant effective frequency delivered over a time period. Energy efficiency and performance optimizations directly impact the resultant effective frequency delivered.

The layout of the IA32\_HWP\_STATUS MSR is shown in Figure 14-9. It provides feedback regarding changes to IA32\_HWP\_CAPABILITIES.Guaranteed\_Performance, IA32\_HWP\_CAPABILITIES.Highest\_Performance, excursions to IA32\_HWP\_CAPABILITIES.Minimum\_Performance, and Peci\_Override entry/exit events. The bit fields are described below:

- **Guaranteed\_Performance\_Change (bit 0, RWCO)** — If set (1), a change to Guaranteed\_Performance has occurred. Software should query IA32\_HWP\_CAPABILITIES.Guaranteed\_Performance value to ascertain the new Guaranteed Performance value and to assess whether to re-adjust HWP hints via IA32\_HWP\_REQUEST. Software must clear this bit by writing a zero (0).
- Bit 1 is reserved and must be zero.
- **Excursion\_To\_Minimum (bit 2, RWCO)** — If set (1), an excursion to Minimum\_Performance of IA32\_HWP\_REQUEST has occurred. Software must clear this bit by writing a zero (0).
- **Highest\_Change (bit 3, RWCO)** — If set (1), a change to Highest Performance has occurred. Software should query IA32\_HWP\_CAPABILITIES to ascertain the new Highest Performance value. Software must clear this bit by writing a zero (0). Interrupts upon Highest Performance change are supported if CPUID[6].EAX[15] is set.
- **PECI\_Override\_Entry (bit 4, RWCO)** — If set (1), an embedded/management controller has started a PECI override of one or more OS control hints (Min, Max, EPP) specified in IA32\_HWP\_REQUEST or IA32\_HWP\_REQUEST\_PKG. Software may query IA32\_HWP\_PECI\_REQUEST\_INFO MSR to ascertain which fields are now overridden via the PECI mechanism and what their values are (see Section 14.4.4.3 for additional details). Software must clear this bit by writing a zero (0). Interrupts upon PECI override entry are supported if CPUID[6].EAX[16] is set.
- **PECI\_Override\_Exit (bit 5, RWCO)** — If set (1), an embedded/management controller has stopped overriding one or more OS control hints (Min, Max, EPP) specified in IA32\_HWP\_REQUEST or IA32\_HWP\_REQUEST\_PKG. Software may query IA32\_HWP\_PECI\_REQUEST\_INFO MSR to ascertain which fields are still overridden via the PECI mechanism and which fields are now back under software control (see Section 14.4.4.3 for additional details). Software must clear this bit by writing a zero (0). Interrupts upon PECI override exit are supported if CPUID[6].EAX[16] is set.
- Bits 63:6 are reserved and must be zero.




Figure 14-9. IA32\_HWP\_STATUS MSR

The status bits of IA32\_HWP\_STATUS must be cleared (0) by software so that a new status condition change will cause the hardware to set the bit again and issue the notification. Status bits are not set for “normal” excursions, e.g., running below Minimum Performance for short durations during C-state exit. Changes to Guaranteed\_Performance, Highest\_Performance, excursions to Minimum\_Performance, or PECI\_Override entry/exit will occur no more than once per second.

The OS can determine the specific reasons for a Guaranteed\_Performance change or an excursion to Minimum\_Performance in IA32\_HWP\_REQUEST by examining the associated status and log bits reported in the IA32\_THERM\_STATUS MSR. The layout of the IA32\_HWP\_STATUS MSR that HWP uses to support software query of HWP feedback is shown in Figure 14-10. The bit fields of IA32\_THERM\_STATUS associated with HWP feedback are described below (Bit fields of IA32\_THERM\_STATUS unrelated to HWP can be found in Section 14.7.5.2).




Figure 14-10. IA32\_THERM\_STATUS Register With HWP Feedback

- Bits 11:0, See Section 14.7.5.2.
- **Current Limit Status (bit 12, RO)** — If set (1), indicates an electrical current limit (e.g. Electrical Design Point/IccMax) is being exceeded and is adversely impacting energy efficiency optimizations.
- **Current Limit Log (bit 13, RWCO)** — If set (1), an electrical current limit has been exceeded that has adversely impacted energy efficiency optimizations since the last clearing of this bit or a reset. This bit is sticky, software may clear this bit by writing a zero (0).
- **Cross-domain Limit Status (bit 14, RO)** — If set (1), indicates another hardware domain (e.g. processor graphics) is currently limiting energy efficiency optimizations in the processor core domain.
- **Cross-domain Limit Log (bit 15, RWCO)** — If set (1), indicates another hardware domain (e.g. processor graphics) has limited energy efficiency optimizations in the processor core domain since the last clearing of this bit or a reset. This bit is sticky, software may clear this bit by writing a zero (0).
- Bits 63:16, See Section 14.7.5.2.

#### 14.4.5.1 Non-Architectural HWP Feedback

The Productive Performance (MSR\_PPERF) MSR (non-architectural) provides hardware's view of workload scalability, which is a rough assessment of the relationship between frequency and workload performance, to software. The layout of the MSR\_PPERF is shown in Figure 14-11.




Figure 14-11. MSR\_PPERF MSR

- **PCNT (bits 63:0, RO)** — Similar to IA32\_APERF but only counts cycles perceived by hardware as contributing to instruction execution (e.g. unhalted and unstalled cycles). This counter increments at the same rate as IA32\_APERF, where the ratio of ( $\Delta PCNT / \Delta ACNT$ ) is an indicator of workload scalability (0% to 100%). Note that values in this register are valid even when HWP is not enabled.

## 14.4.6 HWP Notifications

Processors may support interrupt-based notification of changes to HWP status as indicated by CPUID. If supported, the IA32\_HWP\_INTERRUPT MSR is used to enable interrupt-based notifications. Notification events, when enabled, are delivered using the existing thermal LVT entry. The layout of the IA32\_HWP\_INTERRUPT is shown in Figure 14-12. The bit fields are described below:




Figure 14-12. IA32\_HWP\_INTERRUPT MSR

- **EN\_Guaranteed\_Performance\_Change (bit 0, RW)** — When set (1), an HWP Interrupt will be generated whenever a change to the IA32\_HWP\_CAPABILITIES.Guaranteed\_Performance occurs. The default value is 0 (Interrupt generation is disabled).
- **EN\_Excursion\_Minimum (bit 1, RW)** — When set (1), an HWP Interrupt will be generated whenever the HWP hardware is unable to meet the IA32\_HWP\_REQUEST.Minimum\_Performance setting. The default value is 0 (Interrupt generation is disabled).
- **EN\_Highest\_Change (bit 2, RW)** — When set (1), an HWP Interrupt will be generated whenever a change to the IA32\_HWP\_CAPABILITIES.Highest\_Performance occurs. The default value is 0 (interrupt generation is disabled). Interrupts upon Highest Performance change are supported if CPUID[6].EAX[15] is set.
- **EN\_PECI\_OVERRIDE (bit 3, RW)** — When set (1), an HWP Interrupt will be generated whenever PECI starts or stops overriding any of the three HWP fields described in Section 14.4.4.3. The default value is 0 (interrupt generation is disabled). See Section 14.4.5 and Section 14.4.4.3 for details on how the OS learns what is the current set of HWP fields that are overridden by PECI. Interrupts upon PECI override change are supported if CPUID[6].EAX[16] is set.
- Bits 63:4 are reserved and must be zero.

## 14.4.7 Idle Logical Processor Impact on Core Frequency

Intel processors use one of two schemes for setting core frequency:

1. All cores share same frequency.
2. Each physical core is set to a frequency of its own.

In both cases the two logical processors that share a single physical core are set to the same frequency, so the processor accounts for the IA32\_HWP\_REQUEST MSR fields of both logical processors when defining the core frequency or the whole package frequency.

When CPUID[6].EAX[20] is set and only one logical processor of the two is active, while the other is idle (in any C1 sub-state or in a deeper sleep state), only the active logical processor's IA32\_HWP\_REQUEST MSR fields are considered, i.e., the HWP Request fields of a logical processor in the C1E sub-state or in a deeper sleep state are ignored.

**Note:** when a logical processor is in C1 state its HWP Request fields are accounted for.



## 14.4.8 Fast Write of Uncore MSR (Model Specific Feature)

There are a few logical processor scope MSR's whose values need to be observed outside the logical processor. The WRMSR instruction takes over 1000 cycles to complete (retire) for those MSR's. This overhead forces operating systems to avoid writing them too often whereas in many cases it is preferable that the OS writes them quite frequently for optimal power/performance operation of the processor.

The model specific "Fast Write MSR" feature reduces this overhead by an order of magnitude to a level of 100 cycles for a selected subset of MSR's.

**Note:** Writes to Fast Write MSR's are posted, i.e., when the WRMSR instruction completes, the data may still be "in transit" within the processor. Software can check the status by querying the processor to ensure data is already visible outside the logical processor (see Section 14.4.8.3 for additional details). Once the data is visible outside the logical processor, software is ensured that later writes by the same logical processor to the same MSR will be visible later (will not bypass the earlier writes).

MSR's that are selected for Fast Write are specified in a special capability MSR (see Section 14.4.8.1). Architectural MSR's that existed prior to the introduction of this feature and are selected for Fast Write, thus turning from slow to fast write MSR's, will be noted as such via a new CPUID bit. New MSR's that are fast upon introduction will be documented as such without an additional CPUID bit.

Three model specific MSR's are associated with the feature itself. They enable *enumerating, controlling and monitoring* it. All three are logical processor scope.

### 14.4.8.1 FAST\_UNCORE\_MSRS\_CAPABILITY (Address: 0x65F, Logical Processor Scope)

Operating systems or BIOS can read the FAST\_UNCORE\_MSRS\_CAPABILITY MSR to enumerate those MSR's that are Fast Write MSR's.




Figure 14-13. FAST\_UNCORE\_MSRS\_CAPABILITY MSR

- **FAST\_IA32\_HWP\_REQUEST MSR (bit 0, RO)** — When set (1), indicates that the IA32\_HWP\_REQUEST MSR is supported as a Fast Write MSR. A value of 0 indicates the IA32\_HWP\_REQUEST MSR is not supported as a Fast Write MSR.
- Bits 63:1 are reserved and must be zero.

### 14.4.8.2 FAST\_UNCORE\_MSRS\_CTL (Address: 0x657, Logical Processor Scope)

Operating Systems or BIOS can use the FAST\_UNCORE\_MSRS\_CTL MSR to opt-in or opt-out for fast write of specific MSR's that are enabled for Fast Write by the processor.

**Note:** Not all MSR's that are selected for this feature will necessarily have this opt-in/opt-out option. They may be supported in fast write mode only.




Figure 14-14. FAST\_UNCORE\_MSRS\_CTL MSR

- **FAST\_IA32\_HWP\_REQUEST\_MSR\_ENABLE (bit 0, RW)** — When set (1), enables fast access mode for the IA32\_HWP\_REQUEST MSR and sets the low latency, posted IA32\_HWP\_REQUEST MSR' CPUID[6].EAX[18]. The default value is 0. Note that this bit can only be enabled once from the default value. Once set, writes to this bit are ignored. Only RESET will clear this bit.
- Bits 63:1 are reserved and must be zero.

#### 14.4.8.3 FAST\_UNCORE\_MSRS\_STATUS (Address: 0x65E, Logical Processor Scope)

Software that executes the WRMSR instruction of a Fast Write MSR can check whether the data is already visible outside the logical processor by reading the FAST\_UNCORE\_MSRS\_STATUS MSR. For each Fast Write MSR there is a status bit that indicates whether the data is already visible outside the logical processor or is still in "transit".




Figure 14-15. FAST\_UNCORE\_MSRS\_STATUS MSR

- **FAST\_IA32\_HWP\_REQUEST\_WRITE\_STATUS (bit 0, RO)** — Indicates whether the CPU is still in the middle of writing IA32\_HWP\_REQUEST MSR, even after the WRMSR instruction has retired. A value of 1 indicates the last write of IA32\_HWP\_REQUEST is still ongoing. A value of 0 indicates the last write of IA32\_HWP\_REQUEST is visible outside the logical processor.
- Bits 63:1 are reserved and must be zero.

#### 14.4.9 Fast\_IA32\_HWP\_REQUEST CPUID

IA32\_HWP\_REQUEST is an architectural MSR that exists in processors whose CPUID[6].EAX[7] is set (HWP BASE is enabled). This MSR has logical processor scope, but after its contents are written the contents become visible outside the logical processor. When the FAST\_IA32\_HWP\_REQUEST CPUID[6].EAX[18] bit is set, writes to the IA32\_HWP\_REQUEST MSR are visible outside the logical processor via the "Fast Write" feature described in Section 14.4.8.

#### 14.4.10 Recommendations for OS use of HWP Controls

##### Common Cases of Using HWP

The default HWP control field values are expected to be suitable for many applications. The OS can enable autonomous HWP for these common cases by

- Setting `IA32_HWP_REQUEST.Desired_Performance = 0` (hardware autonomous selection determines the performance target). Set `IA32_HWP_REQUEST.Activity_Window = 0` (enable HW dynamic selection of window size).

To maximize HWP benefit for the common cases, the OS should set

- `IA32_HWP_REQUEST.Minimum_Performance = IA32_HWP_CAPABILITIES.Lowest_Performance` and
- `IA32_HWP_REQUEST.Maximum_Performance = IA32_HWP_CAPABILITIES.Highest_Performance`.

Setting `IA32_HWP_REQUEST.Minimum_Performance = IA32_HWP_REQUEST.Maximum_Performance` is functionally equivalent to using of the `IA32_PERF_CTL` interface and is therefore not recommended (bypassing HWP).

### Calibrating HWP for Application-Specific HWP Optimization

In some applications, the OS may have Quality of Service requirements that may not be met by the default values. The OS can characterize HWP by:

- keeping `IA32_HWP_REQUEST.Minimum_Performance = IA32_HWP_REQUEST.Maximum_Performance` to prevent non-linearity in the characterization process,
- utilizing the range values enumerated from the `IA32_HWP_CAPABILITIES` MSR to program `IA32_HWP_REQUEST` while executing workloads of interest and observing the power and performance result.

The power and performance result of characterization is also influenced by the `IA32_HWP_REQUEST.Energy_Performance_Preference` field, which must also be characterized.

Characterization can be used to set `IA32_HWP_REQUEST.Minimum_Performance` to achieve the required QOS in terms of performance. If `IA32_HWP_REQUEST.Minimum_Performance` is set higher than `IA32_HWP_CAPABILITIES.Guaranteed_Performance` then notification of excursions to Minimum Performance may be continuous.

If autonomous selection does not deliver the required workload performance, the OS should assess the current delivered effective frequency and for the duration of the specific performance requirement set `IA32_HWP_REQUEST.Desired_Performance ≠ 0` and adjust `IA32_HWP_REQUEST.Energy_Performance_Preference` as necessary to achieve the required workload performance. The `MSR_PPERF.PCNT` value can be used to better comprehend the potential performance result from adjustments to `IA32_HWP_REQUEST.Desired_Performance`. The OS should set `IA32_HWP_REQUEST.Desired_Performance = 0` to re-enable autonomous selection.

### Tuning for Maximum Performance or Lowest Power Consumption

Maximum performance will be delivered by setting `IA32_HWP_REQUEST.Minimum_Performance = IA32_HWP_REQUEST.Maximum_Performance = IA32_HWP_CAPABILITIES.Highest_Performance` and setting `IA32_HWP_REQUEST.Energy_Performance_Preference = 0` (performance preference).

Lowest power will be achieved by setting `IA32_HWP_REQUEST.Minimum_Performance = IA32_HWP_REQUEST.Maximum_Performance = IA32_HWP_CAPABILITIES.Lowest_Performance` and setting `IA32_HWP_REQUEST.Energy_Performance_Preference = 0FFH` (energy efficiency preference).

### Mixing Logical Processor and Package Level HWP Field Settings

Using the `IA32_HWP_REQUEST.Package_Control` bit and the five valid bits in that MSR, the OS can mix and match between selecting the Logical Processor scope fields and the Package level fields. For example, the OS can set all logical cores' `IA32_HWP_REQUEST.Package_Control` bit to '1', and for those logical processors if it prefers a different EPP value than the one set in the `IA32_HWP_REQUEST_PKG` MSR, the OS can set the desired EPP value and the EPP valid bit. This overrides the package EPP value for only a subset of the logical processors in the package.

### Additional Guidelines

Set `IA32_HWP_REQUEST.Energy_Performance_Preference` as appropriate for the platform's current mode of operation. For example, a mobile platforms' setting may be towards performance preference when on AC power and more towards energy efficiency when on DC power.

The use of the Running Average Power Limit (RAPL) processor capability (see section 14.7.1) is highly recommended when HWP is enabled. Use of IA32\_HWP\_Request.Maximum\_Performance for thermal control is subject to limitations and can adversely impact the performance of other processor components e.g. Graphics

If default values deliver undesirable performance latency in response to events, the OS should set IA32\_HWP\_REQUEST.Activity\_Window to a low (non-zero) value and IA32\_HWP\_REQUEST.Energy\_Performance\_Preference towards performance (0) for the event duration.

Similarly, for “real-time” threads, set IA32\_HWP\_REQUEST.Energy\_Performance\_Preference towards performance (0) and IA32\_HWP\_REQUEST.Activity\_Window to a low value, e.g. 01H, for the duration of their execution.

When executing low priority work that may otherwise cause the hardware to deliver high performance, set IA32\_HWP\_REQUEST.Activity\_Window to a longer value and reduce the IA32\_HWP\_Request.Maximum\_Performance value as appropriate to control energy efficiency. Adjustments to IA32\_HWP\_REQUEST.Energy\_Performance\_Preference may also be necessary.

## 14.5 HARDWARE DUTY CYCLING (HDC)

Intel processors may contain support for Hardware Duty Cycling (HDC), which enables the processor to autonomously force its components inside the physical package into idle state. For example, the processor may selectively force only the processor cores into an idle state.

HDC is disabled by default on processors that support it. System software can dynamically enable or disable HDC to force one or more components into an idle state or wake up those components previously forced into an idle state. Forced Idling (and waking up) of multiple components in a physical package can be done with one WRMSR to a packaged-scope MSR from any logical processor within the same package.

HDC does not delay events such as timer expiration, but it may affect the latency of short (less than 1 msec) software threads, e.g. if a thread is forced to idle state just before completion and entering a “natural idle”.

HDC forced idle operation can be thought of as operating at a lower effective frequency. The effective average frequency computed by software will include the impact of HDC forced idle.

The primary use of HDC is enable system software to manage low active workloads to increase the package level C6 residency. Additionally, HDC can lower the effective average frequency in case of power or thermal limitation.

When HDC forces a logical processor, a processor core or a physical package to enter an idle state, its C-State is set to C3 or deeper. The deep “C-states” referred to in this section are processor-specific C-states.

### 14.5.1 Hardware Duty Cycling Programming Interfaces

The programming interfaces provided by HDC include the following:

- The CPUID instruction allows software to discover the presence of HDC support in an Intel processor. Specifically, execute CPUID instruction with EAX=06H as input, bit 13 of EAX indicates the processor’s support of the following aspects of HDC.
  - Availability of HDC baseline resource, CPUID.06H:EAX[bit 13]: If this bit is set, HDC provides the following architectural MSRs: IA32\_PKG\_HDC\_CTL, IA32\_PM\_CTL1, and the IA32\_THREAD\_STALL MSRs.
- Additionally, HDC may provide several non-architectural MSR.

**Table 14-2. Architectural and non-Architecture MSRs Related to HDC**

| Address | Architectural | Register Name     | Description   |
|---------|---------------|-------------------|---|
| DB0H    | Y             | IA32_PKG_HDC_CTL  | Package Enable/Disable HDC.   |
| DB1H    | Y             | IA32_PM_CTL1      | Per-logical-processor select control to allow/block HDC forced idling.        |
| DB2H    | Y             | IA32_THREAD_STALL | Accumulate stalled cycles on this logical processor due to HDC forced idling. |

**Table 14-2. Architectural and non-Architecture MSRs Related to HDC**


|      |   |                               |   |
|------|---|-------------------------------|---|
| 653H | N | MSR_CORE_HDC_RESIDENCY        | Core level stalled cycle counter due to HDC forced idling on one or more logical processor.   |
| 655H | N | MSR_PKG_HDC_SHALLOW_RESIDENCY | Accumulate the cycles the package was in C2 <sup>1</sup> state and at least one logical processor was in forced idle  |
| 656H | N | MSR_PKG_HDC_DEEP_RESIDENCY    | Accumulate the cycles the package was in the software specified Cx <sup>1</sup> state and at least one logical processor was in forced idle. Cx is specified in MSR_PKG_HDC_CONFIG_CTL. |
| 652H | N | MSR_PKG_HDC_CONFIG_CTL        | HDC configuration controls  |

**NOTES:**

1. The package “C-states” referred to in this section are processor-specific C-states.

## 14.5.2 Package level Enabling HDC

The layout of the IA32\_PKG\_HDC\_CTL MSR is shown in Figure 14-16. IA32\_PKG\_HDC\_CTL is a writable MSR from any logical processor in a package. The bit fields are described below:

**Figure 14-16. IA32\_PKG\_HDC\_CTL MSR**

- **HDC\_PKG\_Enable (bit 0, R/W)** — Software sets this bit to enable HDC operation by allowing the processor to force to idle all “HDC-allowed” (see Figure 14.5.3) logical processors in the package. Clearing this bit disables HDC operation in the package by waking up all the processor cores that were forced into idle by a previous ‘0’-to-‘1’ transition in IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable. This bit is writable only if CPUID.06H:EAX[bit 13] = 1. Default = zero (0).
- Bits 63:1 are reserved and must be zero.

After processor support is determined via CPUID, system software can enable HDC operation by setting IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable to 1. At reset, IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable is cleared to 0. A ‘0’-to-‘1’ transition in HDC\_PKG\_Enable allows the processor to force to idle all HDC-allowed (indicated by the non-zero state of IA32\_PM\_CTL1[bit 0]) logical processors in the package. A ‘1’-to-‘0’ transition wakes up those HDC force-idled logical processors.

Software can enable or disable HDC using this package level control multiple times from any logical processor in the package. Note the latency of writing a value to the package-visible IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable is longer than the latency of a WRMSR operation to a Logical Processor MSR (as opposed to package level MSR) such as: IA32\_PM\_CTL1 (described in Section 14.5.3). Propagation of the change in IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable and reaching all HDC idled logical processor to be woken up may take on the order of core C6 exit latency.

## 14.5.3 Logical-Processor Level HDC Control

The layout of the IA32\_PM\_CTL1 MSR is shown in Figure 14-17. Each logical processor in a package has its own IA32\_PM\_CTL1 MSR. The bit fields are described below:




Figure 14-17. IA32\_PM\_CTL1 MSR

- **HDC\_Allow\_Block (bit 0, R/W)** — Software sets this bit to allow this logical processors to honor the package-level IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable control. Clearing this bit prevents this logical processor from using the HDC. This bit is writable only if CPUID.06H:EAX[bit 13] = 1. Default = one (1).
- Bits 63:1 are reserved and must be zero.

Fine-grain OS control of HDC operation at the granularity of per-logical-processor is provided by IA32\_PM\_CTL1. At RESET, all logical processors are allowed to participate in HDC operation such that OS can manage HDC using the package-level IA32\_PKG\_HDC\_CTL.

Writes to IA32\_PM\_CTL1 complete with the latency that is typical to WRMSR to a Logical Processor level MSR. When the OS chooses to manage HDC operation at per-logical-processor granularity, it can write to IA32\_PM\_CTL1 on one or more logical processors as desired. Each write to IA32\_PM\_CTL1 must be done by code that executes on the logical processor targeted to be allowed into or blocked from HDC operation.

Blocking one logical processor for HDC operation may have package level impact. For example, the processor may decide to stop duty cycling of all other Logical Processors as well.

The propagation of IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable in a package takes longer than a WRMSR to IA32\_PM\_CTL1. The last completed write to IA32\_PM\_CTL1 on a logical processor will be honored when a '0'-to-'1' transition of IA32\_PKG\_HDC\_CTL.HDC\_PKG\_Enable arrives to a logical processor.

## 14.5.4 HDC Residency Counters

There is a collection of counters available for software to track various residency metrics related to HDC operation. In general, HDC residency time is defined as the time in HDC forced idle state at the granularity of per-logical-processor, per-core, or package. At the granularity of per-core/package-level HDC residency, at least one of the logical processor in a core/package must be in the HDC forced idle state.

### 14.5.4.1 IA32\_THREAD\_STALL

Software can track per-logical-processor HDC residency using the architectural MSR IA32\_THREAD\_STALL. The layout of the IA32\_THREAD\_STALL MSR is shown in Figure 14-18. Each logical processor in a package has its own IA32\_THREAD\_STALL MSR. The bit fields are described below:




Figure 14-18. IA32\_THREAD\_STALL MSR

- **Stall\_Cycle\_Cnt (bits 63:0, R/O)** — Stores accumulated HDC forced-idle cycle count of this processor core since last RESET. This counter increments at the same rate of the TSC. The count is updated only after the logical processor exits from the forced idled C-state. At each update, the number of cycles that the logical processor was stalled due to forced-idle will be added to the counter. This counter is available only if CPUID.06H:EAX[bit 13] = 1. Default = zero (0).

A value of zero in IA32\_THREAD\_STALL indicates either HDC is not supported or the logical processor never serviced any forced HDC idle. A non-zero value in IA32\_THREAD\_STALL indicates the HDC forced-idle residency times of the logical processor. It also indicates the forced-idle cycles due to HDC that could appear as C0 time to traditional OS accounting mechanisms (e.g. time-stamping OS idle/exit events).

Software can read IA32\_THREAD\_STALL irrespective of the state of IA32\_PKG\_HDC\_CTL and IA32\_PM\_CTL1, as long as CPUID.06H:EAX[bit 13] = 1.

#### 14.5.4.2 Non-Architectural HDC Residency Counters

Processors that support HDC operation may provide the following model-specific HDC residency counters.

##### MSR\_CORE\_HDC\_RESIDENCY

Software can track per-core HDC residency using the counter MSR\_CORE\_HDC\_RESIDENCY. This counter increments when the core is in C3 state or deeper (all logical processors in this core are idle due to either HDC or other mechanisms) and at least one of the logical processors is in HDC forced idle state. The layout of the MSR\_CORE\_HDC\_RESIDENCY is shown in Figure 14-19. Each processor core in a package has its own MSR\_CORE\_HDC\_RESIDENCY MSR. The bit fields are described below:




Figure 14-19. MSR\_CORE\_HDC\_RESIDENCY MSR

- **Core\_Cx\_Duty\_Cycle\_Cnt (bits 63:0, R/O)** — Stores accumulated HDC forced-idle cycle count of this processor core since last RESET. This counter increments at the same rate of the TSC. The count is updated only after core C-state exit from a forced idled C-state. At each update, the increment counts cycles when the core is in a Cx state (all its logical processor are idle) and at least one logical processor in this core was forced into idle state due to HDC. If CPUID.06H:EAX[bit 13] = 0, attempt to access this MSR will cause a #GP fault. Default = zero (0).

A value of zero in MSR\_CORE\_HDC\_RESIDENCY indicates either HDC is not supported or this processor core never serviced any forced HDC idle.

##### MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY

The counter MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY allows software to track HDC residency time when the package is in C2 state, all processor cores in the package are not active and at least one logical processor was forced into idle state due to HDC. The layout of the MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY is shown in Figure 14-20. There is one MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY per package. The bit fields are described below:

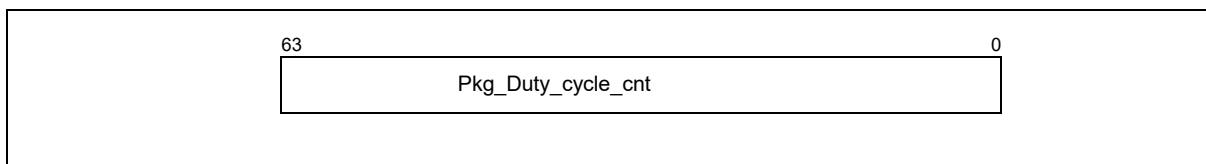


Figure 14-20. MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY MSR

- **Pkg\_Duty\_Cycle\_Cnt (bits 63:0, R/O)** — Stores accumulated HDC forced-idle cycle count of this processor core since last RESET. This counter increments at the same rate of the TSC. Package shallow residency may be implementation specific. In the initial implementation, the threshold is package C2-state. The count is updated only after package C2-state exit from a forced idled C-state. At each update, the increment counts

cycles when the package is in C2 state and at least one processor core in this package was forced into idle state due to HDC. If CPUID.06H:EAX[bit 13] = 0, attempt to access this MSR may cause a #GP fault. Default = zero (0).

A value of zero in MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY indicates either HDC is not supported or this processor package never serviced any forced HDC idle.

### MSR\_PKG\_HDC\_DEEP\_RESIDENCY

The counter MSR\_PKG\_HDC\_DEEP\_RESIDENCY allows software to track HDC residency time when the package is in a software-specified package Cx state, all processor cores in the package are not active and at least one logical processor was forced into idle state due to HDC. Selection of a specific package Cx state can be configured using MSR\_PKG\_HDC\_CONFIG. The layout of the MSR\_PKG\_HDC\_DEEP\_RESIDENCY is shown in Figure 14-21. There is one MSR\_PKG\_HDC\_DEEP\_RESIDENCY per package. The bit fields are described below:




Figure 14-21. MSR\_PKG\_HDC\_DEEP\_RESIDENCY MSR

- **Pkg\_Cx\_Duty\_Cycle\_Cnt (bits 63:0, R/O)** — Stores accumulated HDC forced-idle cycle count of this processor core since last RESET. This counter increments at the same rate of the TSC. The count is updated only after package C-state exit from a forced idle state. At each update, the increment counts cycles when the package is in the software-configured Cx state and at least one processor core in this package was forced into idle state due to HDC. If CPUID.06H:EAX[bit 13] = 0, attempt to access this MSR may cause a #GP fault. Default = zero (0).

A value of zero in MSR\_PKG\_HDC\_SHALLOW\_RESIDENCY indicates either HDC is not supported or this processor package never serviced any forced HDC idle.

### MSR\_PKG\_HDC\_CONFIG

MSR\_PKG\_HDC\_CONFIG allows software to configure the package Cx state that the counter MSR\_PKG\_HDC\_DEEP\_RESIDENCY monitors. The layout of the MSR\_PKG\_HDC\_CONFIG is shown in Figure 14-22. There is one MSR\_PKG\_HDC\_CONFIG per package. The bit fields are described below:




Figure 14-22. MSR\_PKG\_HDC\_CONFIG MSR

- **Pkg\_Cx\_Monitor (bits 2:0, R/W)** — Selects which package C-state the MSR\_HDC\_DEEP\_RESIDENCY counter will monitor. The encoding of the HDC\_Cx\_Monitor field are: 0: no-counting; 1: count package C2 only; 2: count package C3 and deeper; 3: count package C6 and deeper; 4: count package C7 and deeper; other encodings are reserved. If CPUID.06H:EAX[bit 13] = 0, attempt to access this MSR may cause a #GP fault. Default = zero (0).
- Bits 63:3 are reserved and must be zero.



### 14.5.5 MPERF and APERF Counters Under HDC

HDC operation can be thought of as an average effective frequency drop due to all or some of the Logical Processors enter an idle state period.




Figure 14-23. Example of Effective Frequency Reduction and Forced Idle Period of HDC

By default, the IA32\_MPERF counter counts during forced idle periods as if the logical processor was active. The IA32\_APERF counter does not count during forced idle state. This counting convention allows the OS to compute the average effective frequency of the Logical Processor between the last MWAIT exit and the next MWAIT entry (OS visible C0) by  $\Delta\text{ACNT}/\Delta\text{MCNT} * \text{TSC Frequency}$ .

## 14.6 MWAIT EXTENSIONS FOR ADVANCED POWER MANAGEMENT

IA-32 processors may support a number of C-states<sup>1</sup> that reduce power consumption for inactive states. Intel Core Solo and Intel Core Duo processors support both deeper C-state and MWAIT extensions that can be used by OS to implement power management policy.

Software should use CPUID to discover if a target processor supports the enumeration of MWAIT extensions. If CPUID.05H.ECX[Bit 0] = 1, the target processor supports MWAIT extensions and their enumeration (see Chapter 4, "Instruction Set Reference, M-U," of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B*).

If CPUID.05H.ECX[Bit 1] = 1, the target processor supports using interrupts as break-events for MWAIT, even when interrupts are disabled. Use this feature to measure C-state residency as follows:

- Software can write to bit 0 in the MWAIT Extensions register (ECX) when issuing an MWAIT to enter into a processor-specific C-state or sub C-state.
- When a processor comes out of an inactive C-state or sub C-state, software can read a timestamp before an interrupt service routine (ISR) is potentially executed.

CPUID.05H.EDX allows software to enumerate processor-specific C-states and sub C-states available for use with MWAIT extensions. IA-32 processors may support more than one C-state of a given C-state type. These are called sub C-states. Numerically higher C-state have higher power savings and latency (upon entering and exiting) than lower-numbered C-state.

At CPL = 0, system software can specify desired C-state and sub C-state by using the MWAIT hints register (EAX). Processors will not go to C-state and sub C-state deeper than what is specified by the hint register. If CPL > 0 and if MONITOR/MWAIT is supported at CPL > 0, the processor will only enter C1-state (regardless of the C-state request in the hints register).

Executing MWAIT generates an exception on processors operating at a privilege level where MONITOR/MWAIT are not supported.

1. The processor-specific C-states defined in MWAIT extensions can map to ACPI defined C-state types (C0, C1, C2, C3). The mapping relationship depends on the definition of a C-state by processor implementation and is exposed to OSPM by the BIOS using the ACPI defined \_CST table.

**NOTE**

If MWAIT is used to enter a C-state (including sub C-state) that is numerically higher than C1, a store to the address range armed by MONITOR instruction will cause the processor to exit MWAIT if the store was originated by other processor agents. A store from non-processor agent may not cause the processor to exit MWAIT.

## 14.7 THERMAL MONITORING AND PROTECTION

The IA-32 architecture provides the following mechanisms for monitoring temperature and controlling thermal power:


1. The **catastrophic shutdown detector** forces processor execution to stop if the processor's core temperature rises above a preset limit.
2. **Automatic and adaptive thermal monitoring mechanisms** force the processor to reduce its power consumption in order to operate within predetermined temperature limits.
3. The **software controlled clock modulation mechanism** permits operating systems to implement power management policies that reduce power consumption; this is in addition to the reduction offered by automatic thermal monitoring mechanisms.
4. **On-die digital thermal sensor and interrupt mechanisms** permit the OS to manage thermal conditions natively without relying on BIOS or other system board components.

The first mechanism is not visible to software. The other three mechanisms are visible to software using processor feature information returned by executing CPUID with EAX = 1.

The second mechanism includes:

- **Automatic thermal monitoring** provides two modes of operation. One mode modulates the clock duty cycle; the second mode changes the processor's frequency. Both modes are used to control the core temperature of the processor.
- **Adaptive thermal monitoring** can provide flexible thermal management on processors made of multiple cores.

The third mechanism modulates the clock duty cycle of the processor. As shown in Figure 14-24, the phrase 'duty cycle' does not refer to the actual duty cycle of the clock signal. Instead it refers to the time period during which the clock signal is allowed to drive the processor chip. By using the stop clock mechanism to control how often the processor is clocked, processor power consumption can be modulated.



**Figure 14-24. Processor Modulation Through Stop-Clock Mechanism**

For previous automatic thermal monitoring mechanisms, software controlled mechanisms that changed processor operating parameters to impact changes in thermal conditions. Software did not have native access to the native thermal condition of the processor; nor could software alter the trigger condition that initiated software program control.

The fourth mechanism (listed above) provides access to an on-die digital thermal sensor using a model-specific register and uses an interrupt mechanism to alert software to initiate digital thermal monitoring.

## 14.7.1 Catastrophic Shutdown Detector

P6 family processors introduced a thermal sensor that acts as a catastrophic shutdown detector. This catastrophic shutdown detector was also implemented in Pentium 4, Intel Xeon and Pentium M processors. It is always enabled. When processor core temperature reaches a factory preset level, the sensor trips and processor execution is halted until after the next reset cycle.

## 14.7.2 Thermal Monitor

Pentium 4, Intel Xeon and Pentium M processors introduced a second temperature sensor that is factory-calibrated to trip when the processor's core temperature crosses a level corresponding to the recommended thermal design envelop. The trip-temperature of the second sensor is calibrated below the temperature assigned to the catastrophic shutdown detector.

### 14.7.2.1 Thermal Monitor 1

The Pentium 4 processor uses the second temperature sensor in conjunction with a mechanism called Thermal Monitor 1 (TM1) to control the core temperature of the processor. TM1 controls the processor's temperature by modulating the duty cycle of the processor clock. Modulation of duty cycles is processor model specific. Note that the processors STPCLK# pin is not used here; the stop-clock circuitry is controlled internally.

Support for TM1 is indicated by `CPUID.1:EDX.TM[bit 29] = 1`.

TM1 is enabled by setting the thermal-monitor enable flag (bit 3) in `IA32_MISC_ENABLE` [see Chapter 2, "Model-Specific Registers (MSRs)" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*]. Following a power-up or reset, the flag is cleared, disabling TM1. BIOS is required to enable only one automatic thermal monitoring modes. Operating systems and applications must not disable the operation of these mechanisms.

### 14.7.2.2 Thermal Monitor 2

An additional automatic thermal protection mechanism, called Thermal Monitor 2 (TM2), was introduced in the Intel Pentium M processor and also incorporated in newer models of the Pentium 4 processor family. Intel Core Duo and Solo processors, and Intel Core 2 Duo processor family all support TM1 and TM2. TM2 controls the core temperature of the processor by reducing the operating frequency and voltage of the processor and offers a higher performance level for a given level of power reduction than TM1.


TM2 is triggered by the same temperature sensor as TM1. The mechanism to enable TM2 may be implemented differently across various IA-32 processor families with different CPUID signatures in the family encoding value, but will be uniform within an IA-32 processor family.

Support for TM2 is indicated by `CPUID.1:ECX.TM2[bit 8] = 1`.

### 14.7.2.3 Two Methods for Enabling TM2

On processors with CPUID family/model/stepping signature encoded as 0x69n or 0x6Dn (early Pentium M processors), TM2 is enabled if the `TM_SELECT` flag (bit 16) of the `MSR_THERM2_CTL` register is set to 1 (Figure 14-25) and bit 3 of the `IA32_MISC_ENABLE` register is set to 1.


Following a power-up or reset, the `TM_SELECT` flag may be cleared. BIOS is required to enable either TM1 or TM2. Operating systems and applications must not disable mechanisms that enable TM1 or TM2. If bit 3 of the `IA32_MISC_ENABLE` register is set and `TM_SELECT` flag of the `MSR_THERM2_CTL` register is cleared, TM1 is enabled.



**Figure 14-25. MSR\_THERM2\_CTL Register On Processors with CUID Family/Model/Stepping Signature Encoded as 0x69n or 0x6Dn**

On processors introduced after the Pentium 4 processor (this includes most Pentium M processors), the method used to enable TM2 is different. TM2 is enable by setting bit 13 of IA32\_MISC\_ENABLE register to 1. This applies to Intel Core Duo, Core Solo, and Intel Core 2 processor family.

The target operating frequency and voltage for the TM2 transition after TM2 is triggered is specified by the value written to MSR\_THERM2\_CTL, bits 15:0 (Figure 14-26). Following a power-up or reset, BIOS is required to enable at least one of these two thermal monitoring mechanisms. If both TM1 and TM2 are supported, BIOS may choose to enable TM2 instead of TM1. Operating systems and applications must not disable the mechanisms that enable TM1 or TM2; and they must not alter the value in bits 15:0 of the MSR\_THERM2\_CTL register.



**Figure 14-26. MSR\_THERM2\_CTL Register for Supporting TM2**

#### 14.7.2.4 Performance State Transitions and Thermal Monitoring

If the thermal control circuitry (TCC) for thermal monitor (TM1/TM2) is active, writes to the IA32\_PERF\_CTL will effect a new target operating point as follows:

- If TM1 is enabled and the TCC is engaged, the performance state transition can commence before the TCC is disengaged.
- If TM2 is enabled and the TCC is engaged, the performance state transition specified by a write to the IA32\_PERF\_CTL will commence after the TCC has disengaged.

#### 14.7.2.5 Thermal Status Information

The status of the temperature sensor that triggers the thermal monitor (TM1/TM2) is indicated through the thermal status flag and thermal status log flag in the IA32\_THERM\_STATUS MSR (see Figure 14-27).

The functions of these flags are:

- **Thermal Status flag, bit 0** — When set, indicates that the processor core temperature is currently at the trip temperature of the thermal monitor and that the processor power consumption is being reduced via either TM1 or TM2, depending on which is enabled. When clear, the flag indicates that the core temperature is below the thermal monitor trip temperature. This flag is read only.
- **Thermal Status Log flag, bit 1** — When set, indicates that the thermal sensor has tripped since the last power-up or reset or since the last time that software cleared this flag. This flag is a sticky bit; once set it remains set until cleared by software or until a power-up or reset of the processor. The default state is clear.




Figure 14-27. IA32\_THERM\_STATUS MSR

After the second temperature sensor has been tripped, the thermal monitor (TM1/TM2) will remain engaged for a minimum time period (on the order of 1 ms). The thermal monitor will remain engaged until the processor core temperature drops below the preset trip temperature of the temperature sensor, taking hysteresis into account.

While the processor is in a stop-clock state, interrupts will be blocked from interrupting the processor. This holding off of interrupts increases the interrupt latency, but does not cause interrupts to be lost. Outstanding interrupts remain pending until clock modulation is complete.

The thermal monitor can be programmed to generate an interrupt to the processor when the thermal sensor is tripped. The delivery mode, mask and vector for this interrupt can be programmed through the thermal entry in the local APIC's LVT (see Section 10.5.1, "Local Vector Table"). The low-temperature interrupt enable and high-temperature interrupt enable flags in the IA32\_THERM\_INTERRUPT MSR (see Figure 14-28) control when the interrupt is generated; that is, on a transition from a temperature below the trip point to above and/or vice-versa.




Figure 14-28. IA32\_THERM\_INTERRUPT MSR

- **High-Temperature Interrupt Enable flag, bit 0** — Enables an interrupt to be generated on the transition from a low-temperature to a high-temperature when set; disables the interrupt when clear. (R/W).
- **Low-Temperature Interrupt Enable flag, bit 1** — Enables an interrupt to be generated on the transition from a high-temperature to a low-temperature when set; disables the interrupt when clear.

The thermal monitor interrupt can be masked by the thermal LVT entry. After a power-up or reset, the low-temperature interrupt enable and high-temperature interrupt enable flags in the IA32\_THERM\_INTERRUPT MSR are cleared (interrupts are disabled) and the thermal LVT entry is set to mask interrupts. This interrupt should be handled either by the operating system or system management mode (SMM) code.

Note that the operation of the thermal monitoring mechanism has no effect upon the clock rate of the processor's internal high-resolution timer (time stamp counter).

### 14.7.2.6 Adaptive Thermal Monitor

The Intel Core 2 Duo processor family supports enhanced thermal management mechanism, referred to as Adaptive Thermal Monitor (Adaptive TM).

Unlike TM2, Adaptive TM is not limited to one TM2 transition target. During a thermal trip event, Adaptive TM (if enabled) selects an optimal target operating point based on whether or not the current operating point has effectively cooled the processor.

Similar to TM2, Adaptive TM is enable by BIOS. The BIOS is required to test the TM1 and TM2 feature flags and enable all available thermal control mechanisms (including Adaptive TM) at platform initiation.

Adaptive TM is available only to a subset of processors that support TM2.

In each chip-multiprocessing (CMP) silicon die, each core has a unique thermal sensor that triggers independently. These thermal sensor can trigger TM1 or TM2 transitions in the same manner as described in Section 14.7.2.1 and Section 14.7.2.2. The trip point of the thermal sensor is not programmable by software since it is set during the fabrication of the processor.

Each thermal sensor in a processor core may be triggered independently to engage thermal management features. In Adaptive TM, both cores will transition to a lower frequency and/or lower voltage level if one sensor is triggered. Triggering of this sensor is visible to software via the thermal interrupt LVT entry in the local APIC of a given core.

### 14.7.3 Software Controlled Clock Modulation

Pentium 4, Intel Xeon and Pentium M processors also support software-controlled clock modulation. This provides a means for operating systems to implement a power management policy to reduce the power consumption of the processor. Here, the stop-clock duty cycle is controlled by software through the IA32\_CLOCK\_MODULATION MSR (see Figure 14-29).




Figure 14-29. IA32\_CLOCK\_MODULATION MSR

The IA32\_CLOCK\_MODULATION MSR contains the following flag and field used to enable software-controlled clock modulation and to select the clock modulation duty cycle:

- **On-Demand Clock Modulation Enable, bit 4** — Enables on-demand software controlled clock modulation when set; disables software-controlled clock modulation when clear.
- **On-Demand Clock Modulation Duty Cycle, bits 1 through 3** — Selects the on-demand clock modulation duty cycle (see Table 14-3). This field is only active when the on-demand clock modulation enable flag is set.

Note that the on-demand clock modulation mechanism (like the thermal monitor) controls the processor's stop-clock circuitry internally to modulate the clock signal. The STPCLK# pin is not used in this mechanism.

Table 14-3. On-Demand Clock Modulation Duty Cycle Field Encoding

| Duty Cycle Field Encoding | Duty Cycle      |
|---------------------------|-----------------|
| 000B                      | Reserved        |
| 001B                      | 12.5% (Default) |
| 010B                      | 25.0%           |
| 011B                      | 37.5%           |
| 100B                      | 50.0%           |
| 101B                      | 63.5%           |
| 110B                      | 75%             |
| 111B                      | 87.5%           |

The on-demand clock modulation mechanism can be used to control processor power consumption. Power management software can write to the IA32\_CLOCK\_MODULATION MSR to enable clock modulation and to select a modulation duty cycle. If on-demand clock modulation and TM1 are both enabled and the thermal status of the processor is hot (bit 0 of the IA32\_THERM\_STATUS MSR is set), clock modulation at the duty cycle specified by TM1 takes precedence, regardless of the setting of the on-demand clock modulation duty cycle.

For Hyper-Threading Technology enabled processors, the IA32\_CLOCK\_MODULATION register is duplicated for each logical processor. In order for the On-demand clock modulation feature to work properly, the feature must be enabled on all the logical processors within a physical processor. If the programmed duty cycle is not identical for all the logical processors, the processor core clock will modulate to the highest duty cycle programmed for processors with any of the following CPUID DisplayFamily\_DisplayModel signatures (see CPUID instruction in Chapter 3, “Instruction Set Reference, A-L” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2A*): 06\_1A, 06\_1C, 06\_1E, 06\_1F, 06\_25, 06\_26, 06\_27, 06\_2C, 06\_2E, 06\_2F, 06\_35, 06\_36, and 0F\_xx. For all other processors, if the programmed duty cycle is not identical for all logical processors in the same core, the processor core will modulate at the lowest programmed duty cycle.

For multiple processor cores in a physical package, each processor core can modulate to a programmed duty cycle independently.

For the P6 family processors, on-demand clock modulation was implemented through the chipset, which controlled clock modulation through the processor’s STPCLK# pin.

### 14.7.3.1 Extension of Software Controlled Clock Modulation

Extension of the software controlled clock modulation facility supports on-demand clock modulation duty cycle with 4-bit dynamic range (increased from 3-bit range). Granularity of clock modulation duty cycle is increased to 6.25% (compared to 12.5%).

Four bit dynamic range control is provided by using bit 0 in conjunction with bits 3:1 of the IA32\_CLOCK\_MODULATION MSR (see Figure 14-30).




Figure 14-30. IA32\_CLOCK\_MODULATION MSR with Clock Modulation Extension

Extension to software controlled clock modulation is supported only if CPUID.06H:EAX[Bit 5] = 1. If CPUID.06H:EAX[Bit 5] = 0, then bit 0 of IA32\_CLOCK\_MODULATION is reserved.

## 14.7.4 Detection of Thermal Monitor and Software Controlled Clock Modulation Facilities

The ACPI flag (bit 22) of the CPUID feature flags indicates the presence of the IA32\_THERM\_STATUS, IA32\_THERM\_INTERRUPT, IA32\_CLOCK\_MODULATION MSRs, and the xAPIC thermal LVT entry.

The TM1 flag (bit 29) of the CPUID feature flags indicates the presence of the automatic thermal monitoring facilities that modulate clock duty cycles.

### 14.7.4.1 Detection of Software Controlled Clock Modulation Extension

Processor’s support of software controlled clock modulation extension is indicated by CPUID.06H:EAX[Bit 5] = 1.

## 14.7.5 On Die Digital Thermal Sensors

On die digital thermal sensor can be read using an MSR (no I/O interface). In Intel Core Duo processors, each core has a unique digital sensor whose temperature is accessible using an MSR. The digital thermal sensor is the preferred method for reading the die temperature because (a) it is located closer to the hottest portions of the die, (b) it enables software to accurately track the die temperature and the potential activation of thermal throttling.

### 14.7.5.1 Digital Thermal Sensor Enumeration

The processor supports a digital thermal sensor if CPUID.06H.EAX[0] = 1. If the processor supports digital thermal sensor, EBX[bits 3:0] determine the number of thermal thresholds that are available for use.

Software sets thermal thresholds by using the IA32\_THERM\_INTERRUPT MSR. Software reads output of the digital thermal sensor using the IA32\_THERM\_STATUS MSR.

### 14.7.5.2 Reading the Digital Sensor

Unlike traditional analog thermal devices, the output of the digital thermal sensor is a temperature relative to the maximum supported operating temperature of the processor.

Temperature measurements returned by digital thermal sensors are always at or below TCC activation temperature. Critical temperature conditions are detected using the “Critical Temperature Status” bit. When this bit is set, the processor is operating at a critical temperature and immediate shutdown of the system should occur. Once the “Critical Temperature Status” bit is set, reliable operation is not guaranteed.

See Figure 14-31 for the layout of IA32\_THERM\_STATUS MSR. Bit fields include:

- **Thermal Status (bit 0, RO)** — This bit indicates whether the digital thermal sensor high-temperature output signal (PROCHOT#) is currently active. Bit 0 = 1 indicates the feature is active. This bit may not be written by software; it reflects the state of the digital thermal sensor.
- **Thermal Status Log (bit 1, R/WCO)** — This is a sticky bit that indicates the history of the thermal sensor high temperature output signal (PROCHOT#). Bit 1 = 1 if PROCHOT# has been asserted since a previous RESET or the last time software cleared the bit. Software may clear this bit by writing a zero.
- **PROCHOT# or FORCEPR# Event (bit 2, RO)** — Indicates whether PROCHOT# or FORCEPR# is being asserted by another agent on the platform.




Figure 14-31. IA32\_THERM\_STATUS Register

- **PROCHOT# or FORCEPR# Log (bit 3, R/WCO)** — Sticky bit that indicates whether PROCHOT# or FORCEPR# has been asserted by another agent on the platform since the last clearing of this bit or a reset. If bit 3 = 1, PROCHOT# or FORCEPR# has been externally asserted. Software may clear this bit by writing a zero. External PROCHOT# assertions are only acknowledged if the Bidirectional Prochot feature is enabled.
- **Critical Temperature Status (bit 4, RO)** — Indicates whether the critical temperature detector output signal is currently active. If bit 4 = 1, the critical temperature detector output signal is currently active.



- **Critical Temperature Log (bit 5, R/WCO)** — Sticky bit that indicates whether the critical temperature detector output signal has been asserted since the last clearing of this bit or reset. If bit 5 = 1, the output signal has been asserted. Software may clear this bit by writing a zero.
- **Thermal Threshold #1 Status (bit 6, RO)** — Indicates whether the actual temperature is currently higher than or equal to the value set in Thermal Threshold #1. If bit 6 = 0, the actual temperature is lower. If bit 6 = 1, the actual temperature is greater than or equal to TT#1. Quantitative information of actual temperature can be inferred from Digital Readout, bits 22:16.
- **Thermal Threshold #1 Log (bit 7, R/WCO)** — Sticky bit that indicates whether the Thermal Threshold #1 has been reached since the last clearing of this bit or a reset. If bit 7 = 1, the Threshold #1 has been reached. Software may clear this bit by writing a zero.
- **Thermal Threshold #2 Status (bit 8, RO)** — Indicates whether actual temperature is currently higher than or equal to the value set in Thermal Threshold #2. If bit 8 = 0, the actual temperature is lower. If bit 8 = 1, the actual temperature is greater than or equal to TT#2. Quantitative information of actual temperature can be inferred from Digital Readout, bits 22:16.
- **Thermal Threshold #2 Log (bit 9, R/WCO)** — Sticky bit that indicates whether the Thermal Threshold #2 has been reached since the last clearing of this bit or a reset. If bit 9 = 1, the Thermal Threshold #2 has been reached. Software may clear this bit by writing a zero.
- **Power Limitation Status (bit 10, RO)** — Indicates whether the processor is currently operating below OS-requested P-state (specified in IA32\_PERF\_CTL) or OS-requested clock modulation duty cycle (specified in IA32\_CLOCK\_MODULATION). This field is supported only if CPUID.06H:EAX[bit 4] = 1. Package level power limit notification can be delivered independently to IA32\_PACKAGE\_THERM\_STATUS MSR.
- **Power Notification Log (bit 11, R/WCO)** — Sticky bit that indicates the processor went below OS-requested P-state or OS-requested clock modulation duty cycle since the last clearing of this or RESET. This field is supported only if CPUID.06H:EAX[bit 4] = 1. Package level power limit notification is indicated independently in IA32\_PACKAGE\_THERM\_STATUS MSR.
- **Digital Readout (bits 22:16, RO)** — Digital temperature reading in 1 degree Celsius relative to the TCC activation temperature.
  - 0: TCC Activation temperature,
  - 1: (TCC Activation - 1) , etc. See the processor's data sheet for details regarding TCC activation.
 A lower reading in the Digital Readout field (bits 22:16) indicates a higher actual temperature.
- **Resolution in Degrees Celsius (bits 30:27, RO)** — Specifies the resolution (or tolerance) of the digital thermal sensor. The value is in degrees Celsius. It is recommended that new threshold values be offset from the current temperature by at least the resolution + 1 in order to avoid hysteresis of interrupt generation.
- **Reading Valid (bit 31, RO)** — Indicates if the digital readout in bits 22:16 is valid. The readout is valid if bit 31 = 1.

Changes to temperature can be detected using two thresholds (see Figure 14-32); one is set above and the other below the current temperature. These thresholds have the capability of generating interrupts using the core's local APIC which software must then service. Note that the local APIC entries used by these thresholds are also used by the Intel® Thermal Monitor; it is up to software to determine the source of a specific interrupt.




Figure 14-32. IA32\_THERM\_INTERRUPT Register

See Figure 14-32 for the layout of IA32\_THERM\_INTERRUPT MSR. Bit fields include:

- **High-Temperature Interrupt Enable (bit 0, R/W)** — This bit allows the BIOS to enable the generation of an interrupt on the transition from low-temperature to a high-temperature threshold. Bit 0 = 0 (default) disables interrupts; bit 0 = 1 enables interrupts.
- **Low-Temperature Interrupt Enable (bit 1, R/W)** — This bit allows the BIOS to enable the generation of an interrupt on the transition from high-temperature to a low-temperature (TCC de-activation). Bit 1 = 0 (default) disables interrupts; bit 1 = 1 enables interrupts.
- **PROCHOT# Interrupt Enable (bit 2, R/W)** — This bit allows the BIOS or OS to enable the generation of an interrupt when PROCHOT# has been asserted by another agent on the platform and the Bidirectional Prochot feature is enabled. Bit 2 = 0 disables the interrupt; bit 2 = 1 enables the interrupt.
- **FORCEPR# Interrupt Enable (bit 3, R/W)** — This bit allows the BIOS or OS to enable the generation of an interrupt when FORCEPR# has been asserted by another agent on the platform. Bit 3 = 0 disables the interrupt; bit 3 = 1 enables the interrupt.
- **Critical Temperature Interrupt Enable (bit 4, R/W)** — Enables the generation of an interrupt when the Critical Temperature Detector has detected a critical thermal condition. The recommended response to this condition is a system shutdown. Bit 4 = 0 disables the interrupt; bit 4 = 1 enables the interrupt.
- **Threshold #1 Value (bits 14:8, R/W)** — A temperature threshold, encoded relative to the TCC Activation temperature (using the same format as the Digital Readout). This threshold is compared against the Digital Readout and is used to generate the Thermal Threshold #1 Status and Log bits as well as the Threshold #1 thermal interrupt delivery.
- **Threshold #1 Interrupt Enable (bit 15, R/W)** — Enables the generation of an interrupt when the actual temperature crosses the Threshold #1 setting in any direction. Bit 15 = 1 enables the interrupt; bit 15 = 0 disables the interrupt.
- **Threshold #2 Value (bits 22:16, R/W)** — A temperature threshold, encoded relative to the TCC Activation temperature (using the same format as the Digital Readout). This threshold is compared against the Digital Readout and is used to generate the Thermal Threshold #2 Status and Log bits as well as the Threshold #2 thermal interrupt delivery.
- **Threshold #2 Interrupt Enable (bit 23, R/W)** — Enables the generation of an interrupt when the actual temperature crosses the Threshold #2 setting in any direction. Bit 23 = 1 enables the interrupt; bit 23 = 0 disables the interrupt.
- **Power Limit Notification Enable (bit 24, R/W)** — Enables the generation of power notification events when the processor went below OS-requested P-state or OS-requested clock modulation duty cycle. This field is supported only if CPUID.06H:EAX[bit 4] = 1. Package level power limit notification can be enabled independently by IA32\_PACKAGE\_THERM\_INTERRUPT MSR.

### 14.7.6 Power Limit Notification

Platform firmware may be capable of specifying a power limit to restrict power delivered to a platform component, such as a physical processor package. This constraint imposed by platform firmware may occasionally cause the processor to operate below OS-requested P or T-state. A power limit notification event can be delivered using the existing thermal LVT entry in the local APIC.

Software can enumerate the presence of the processor’s support for power limit notification by verifying CPUID.06H:EAX[bit 4] = 1.

If CPUID.06H:EAX[bit 4] = 1, then IA32\_THERM\_INTERRUPT and IA32\_THERM\_STATUS provides the following facility to manage power limit notification:

- Bits 10 and 11 in IA32\_THERM\_STATUS informs software of the occurrence of processor operating below OS-requested P-state or clock modulation duty cycle setting (see Figure 14-31).
- Bit 24 in IA32\_THERM\_INTERRUPT enables the local APIC to deliver a thermal event when the processor went below OS-requested P-state or clock modulation duty cycle setting (see Figure 14-32).

## 14.8 PACKAGE LEVEL THERMAL MANAGEMENT

The thermal management facilities like IA32\_THERM\_INTERRUPT and IA32\_THERM\_STATUS are often implemented with a processor core granularity. To facilitate software manage thermal events from a package level granularity, two architectural MSR is provided for package level thermal management. The IA32\_PACKAGE\_THERM\_STATUS and IA32\_PACKAGE\_THERM\_INTERRUPT MSRs use similar interfaces as IA32\_THERM\_STATUS and IA32\_THERM\_INTERRUPT, but are shared in each physical processor package.

Software can enumerate the presence of the processor’s support for package level thermal management facility (IA32\_PACKAGE\_THERM\_STATUS and IA32\_PACKAGE\_THERM\_INTERRUPT) by verifying CPUID.06H:EAX[bit 6] = 1.

The layout of IA32\_PACKAGE\_THERM\_STATUS MSR is shown in Figure 14-33.




Figure 14-33. IA32\_PACKAGE\_THERM\_STATUS Register

- **Package Thermal Status (bit 0, RO)** — This bit indicates whether the digital thermal sensor high-temperature output signal (PROCHOT#) for the package is currently active. Bit 0 = 1 indicates the feature is active. This bit may not be written by software; it reflects the state of the digital thermal sensor.

- **Package Thermal Status Log (bit 1, R/WCO)** — This is a sticky bit that indicates the history of the thermal sensor high temperature output signal (PROCHOT#) of the package. Bit 1 = 1 if package PROCHOT# has been asserted since a previous RESET or the last time software cleared the bit. Software may clear this bit by writing a zero.
- **Package PROCHOT# Event (bit 2, RO)** — Indicates whether package PROCHOT# is being asserted by another agent on the platform.
- **Package PROCHOT# Log (bit 3, R/WCO)** — Sticky bit that indicates whether package PROCHOT# has been asserted by another agent on the platform since the last clearing of this bit or a reset. If bit 3 = 1, package PROCHOT# has been externally asserted. Software may clear this bit by writing a zero.
- **Package Critical Temperature Status (bit 4, RO)** — Indicates whether the package critical temperature detector output signal is currently active. If bit 4 = 1, the package critical temperature detector output signal is currently active.
- **Package Critical Temperature Log (bit 5, R/WCO)** — Sticky bit that indicates whether the package critical temperature detector output signal has been asserted since the last clearing of this bit or reset. If bit 5 = 1, the output signal has been asserted. Software may clear this bit by writing a zero.
- **Package Thermal Threshold #1 Status (bit 6, RO)** — Indicates whether the actual package temperature is currently higher than or equal to the value set in Package Thermal Threshold #1. If bit 6 = 0, the actual temperature is lower. If bit 6 = 1, the actual temperature is greater than or equal to PTT#1. Quantitative information of actual package temperature can be inferred from Package Digital Readout, bits 22:16.
- **Package Thermal Threshold #1 Log (bit 7, R/WCO)** — Sticky bit that indicates whether the Package Thermal Threshold #1 has been reached since the last clearing of this bit or a reset. If bit 7 = 1, the Package Thermal Threshold #1 has been reached. Software may clear this bit by writing a zero.
- **Package Thermal Threshold #2 Status (bit 8, RO)** — Indicates whether actual package temperature is currently higher than or equal to the value set in Package Thermal Threshold #2. If bit 8 = 0, the actual temperature is lower. If bit 8 = 1, the actual temperature is greater than or equal to PTT#2. Quantitative information of actual temperature can be inferred from Package Digital Readout, bits 22:16.
- **Package Thermal Threshold #2 Log (bit 9, R/WCO)** — Sticky bit that indicates whether the Package Thermal Threshold #2 has been reached since the last clearing of this bit or a reset. If bit 9 = 1, the Package Thermal Threshold #2 has been reached. Software may clear this bit by writing a zero.
- **Package Power Limitation Status (bit 10, RO)** — Indicates package power limit is forcing one or more processors to operate below OS-requested P-state. Note that package power limit violation may be caused by processor cores or by devices residing in the uncore. Software can examine IA32\_THERM\_STATUS to determine if the cause originates from a processor core (see Figure 14-31).
- **Package Power Notification Log (bit 11, R/WCO)** — Sticky bit that indicates any processor in the package went below OS-requested P-state or OS-requested clock modulation duty cycle since the last clearing of this or RESET.
- **Package Digital Readout (bits 22:16, RO)** — Package digital temperature reading in 1 degree Celsius relative to the package TCC activation temperature.
  - 0: Package TCC Activation temperature,
  - 1: (PTCC Activation - 1) , etc. See the processor's data sheet for details regarding PTCC activation.
 A lower reading in the Package Digital Readout field (bits 22:16) indicates a higher actual temperature.

The layout of IA32\_PACKAGE\_THERM\_INTERRUPT MSR is shown in Figure 14-34.




Figure 14-34. IA32\_PACKAGE\_THERM\_INTERRUPT Register

- **Package High-Temperature Interrupt Enable (bit 0, R/W)** — This bit allows the BIOS to enable the generation of an interrupt on the transition from low-temperature to a package high-temperature threshold. Bit 0 = 0 (default) disables interrupts; bit 0 = 1 enables interrupts.
- **Package Low-Temperature Interrupt Enable (bit 1, R/W)** — This bit allows the BIOS to enable the generation of an interrupt on the transition from high-temperature to a low-temperature (TCC de-activation). Bit 1 = 0 (default) disables interrupts; bit 1 = 1 enables interrupts.
- **Package PROCHOT# Interrupt Enable (bit 2, R/W)** — This bit allows the BIOS or OS to enable the generation of an interrupt when Package PROCHOT# has been asserted by another agent on the platform and the Bidirectional Prochot feature is enabled. Bit 2 = 0 disables the interrupt; bit 2 = 1 enables the interrupt.
- **Package Critical Temperature Interrupt Enable (bit 4, R/W)** — Enables the generation of an interrupt when the Package Critical Temperature Detector has detected a critical thermal condition. The recommended response to this condition is a system shutdown. Bit 4 = 0 disables the interrupt; bit 4 = 1 enables the interrupt.
- **Package Threshold #1 Value (bits 14:8, R/W)** — A temperature threshold, encoded relative to the Package TCC Activation temperature (using the same format as the Digital Readout). This threshold is compared against the Package Digital Readout and is used to generate the Package Thermal Threshold #1 Status and Log bits as well as the Package Threshold #1 thermal interrupt delivery.
- **Package Threshold #1 Interrupt Enable (bit 15, R/W)** — Enables the generation of an interrupt when the actual temperature crosses the Package Threshold #1 setting in any direction. Bit 15 = 1 enables the interrupt; bit 15 = 0 disables the interrupt.
- **Package Threshold #2 Value (bits 22:16, R/W)** —A temperature threshold, encoded relative to the PTCC Activation temperature (using the same format as the Package Digital Readout). This threshold is compared against the Package Digital Readout and is used to generate the Package Thermal Threshold #2 Status and Log bits as well as the Package Threshold #2 thermal interrupt delivery.
- **Package Threshold #2 Interrupt Enable (bit 23, R/W)** — Enables the generation of an interrupt when the actual temperature crosses the Package Threshold #2 setting in any direction. Bit 23 = 1 enables the interrupt; bit 23 = 0 disables the interrupt.
- **Package Power Limit Notification Enable (bit 24, R/W)** — Enables the generation of package power notification events.

### 14.8.1 Support for Passive and Active cooling

Passive and active cooling may be controlled by the OS power management agent through ACPI control methods. On platforms providing package level thermal management facility described in the previous section, it is recommended that active cooling (FAN control) should be driven by measuring the package temperature using the IA32\_PACKAGE\_THERM\_INTERRUPT MSR.

Passive cooling (frequency throttling) should be driven by measuring (a) the core and package temperatures, or (b) only the package temperature. If measured package temperature led the power management agent to choose which core to execute passive cooling, then all cores need to execute passive cooling. Core temperature is measured using the IA32\_THERMAL\_STATUS and IA32\_THERMAL\_INTERRUPT MSRs. The exact implementation details depend on the platform firmware and possible solutions include defining two different thermal zones (one for core temperature and passive cooling and the other for package temperature and active cooling).

## 14.9 PLATFORM SPECIFIC POWER MANAGEMENT SUPPORT

This section covers power management interfaces that are not architectural but addresses the power management needs of several platform specific components. Specifically, RAPL (Running Average Power Limit) interfaces provide mechanisms to enforce power consumption limit. Power limiting usages have specific usages in client and server platforms.

For client platform power limit control and for server platforms used in a data center, the following power and thermal related usages are desirable:

- Platform Thermal Management: Robust mechanisms to manage component, platform, and group-level thermals, either proactively or reactively (e.g., in response to a platform-level thermal trip point).
- Platform Power Limiting: More deterministic control over the system's power consumption, for example to meet battery life targets on rack-level or container-level power consumption goals within a datacenter.
- Power/Performance Budgeting: Efficient means to control the power consumed (and therefore the sustained performance delivered) within and across platforms.

The server and client usage models are addressed by RAPL interfaces, which expose multiple domains of power rationing within each processor socket. Generally, these RAPL domains may be viewed to include hierarchically:

- Package domain is the processor die.
- Memory domain includes the directly-attached DRAM; an additional power plane may constitute a separate domain.

In order to manage the power consumed across multiple sockets via RAPL, individual limits must be programmed for each processor complex. Programming specific RAPL domain across multiple sockets is not supported.

### 14.9.1 RAPL Interfaces

RAPL interfaces consist of non-architectural MSRs. Each RAPL domain supports the following set of capabilities, some of which are optional as stated below.

- Power limit - MSR interfaces to specify power limit, time window; lock bit, clamp bit etc.
- Energy Status - Power metering interface providing energy consumption information.
- Perf Status (Optional) - Interface providing information on the performance effects (regression) due to power limits. It is defined as a duration metric that measures the power limit effect in the respective domain. The meaning of duration is domain specific.
- Power Info (Optional) - Interface providing information on the range of parameters for a given domain, minimum power, maximum power etc.
- Policy (Optional) - 4-bit priority information that is a hint to hardware for dividing budget between sub-domains in a parent domain.

Each of the above capabilities requires specific units in order to describe them. Power is expressed in Watts, Time is expressed in Seconds, and Energy is expressed in Joules. Scaling factors are supplied to each unit to make the information presented meaningful in a finite number of bits. Units for power, energy, and time are exposed in the read-only MSR\_RAPL\_POWER\_UNIT MSR.




Figure 14-35. MSR\_RAPL\_POWER\_UNIT Register

MSR\_RAPL\_POWER\_UNIT (Figure 14-35) provides the following information across all RAPL domains:

- **Power Units (bits 3:0):** Power related information (in Watts) is based on the multiplier,  $1/2^{PU}$ ; where PU is an unsigned integer represented by bits 3:0. Default value is 0011b, indicating power unit is in 1/8 Watts increment.
- **Energy Status Units (bits 12:8):** Energy related information (in Joules) is based on the multiplier,  $1/2^{ESU}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 10000b, indicating energy status unit is in 15.3 micro-Joules increment.
- **Time Units (bits 19:16):** Time related information (in Seconds) is based on the multiplier,  $1/2^{TU}$ ; where TU is an unsigned integer represented by bits 19:16. Default value is 1010b, indicating time unit is in 976 micro-seconds increment.

### 14.9.2 RAPL Domains and Platform Specificity

The specific RAPL domains available in a platform vary across product segments. Platforms targeting the client segment support the following RAPL domain hierarchy:

- Package
- Two power planes: PPO and PP1 (PP1 may reflect to uncore devices)

Platforms targeting the server segment support the following RAPL domain hierarchy:

- Package
- Power plane: PPO
- DRAM

Each level of the RAPL hierarchy provides a respective set of RAPL interface MSRs. Table 14-4 lists the RAPL MSR interfaces available for each RAPL domain. The power limit MSR of each RAPL domain is located at offset 0 relative to an MSR base address which is non-architectural (see Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*). The energy status MSR of each domain is located at offset 1 relative to the MSR base address of respective domain.

Table 14-4. RAPL MSR Interfaces and RAPL Domains

| Domain | Power Limit (Offset 0) | Energy Status (Offset 1) | Policy (Offset 2) | Perf Status (Offset 3) | Power Info (Offset 4) |
|--------|------------------------|--------------------------|-------------------|------------------------|-----------------------|
| PKG    | MSR_PKG_POWER_LIMIT    | MSR_PKG_ENERGY_STATUS    | RESERVED          | MSR_PKG_PERF_STATUS    | MSR_PKG_POWER_INFO    |
| DRAM   | MSR_DRAM_POWER_LIMIT   | MSR_DRAM_ENERGY_STATUS   | RESERVED          | MSR_DRAM_PERF_STATUS   | MSR_DRAM_POWER_INFO   |
| PPO    | MSR_PPO_POWER_LIMIT    | MSR_PPO_ENERGY_STATUS    | MSR_PPO_POLICY    | MSR_PPO_PERF_STATUS    | RESERVED              |

**Table 14-4. RAPL MSR Interfaces and RAPL Domains**

|     |                     |                       |                |          |          |
|-----|---------------------|-----------------------|----------------|----------|----------|
| PP1 | MSR_PP1_POWER_LIMIT | MSR_PP1_ENERGY_STATUS | MSR_PP1_POLICY | RESERVED | RESERVED |
|-----|---------------------|-----------------------|----------------|----------|----------|

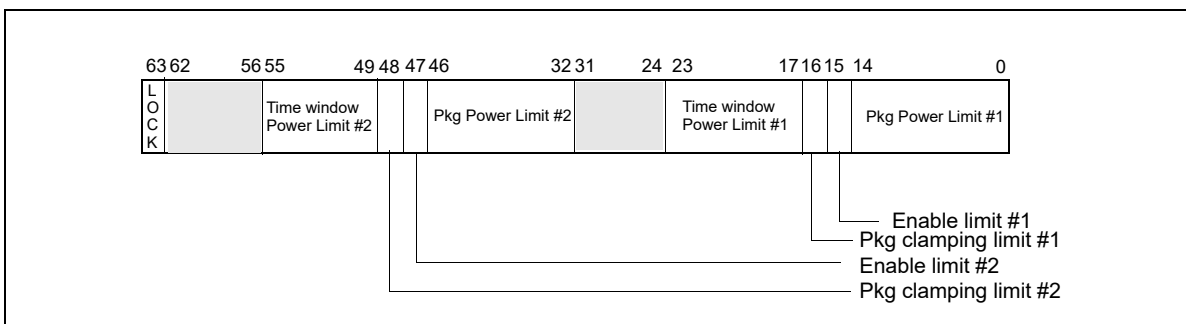
The presence of the optional MSR interfaces (the three right-most columns of Table 14-4) may be model-specific. See Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4* for details.

### 14.9.3 Package RAPL Domain

The MSR interfaces defined for the package RAPL domain are:

- MSR\_PKG\_POWER\_LIMIT allows software to set power limits for the package and measurement attributes associated with each limit,
- MSR\_PKG\_ENERGY\_STATUS reports measured actual energy usage,
- MSR\_PKG\_POWER\_INFO reports the package power range information for RAPL usage.

MSR\_PKG\_PERF\_STATUS can report the performance impact of power limiting, but its availability may be model-specific.



**Figure 14-36. MSR\_PKG\_POWER\_LIMIT Register**

MSR\_PKG\_POWER\_LIMIT allows a software agent to define power limitation for the package domain. Power limitation is defined in terms of average power usage (Watts) over a time window specified in MSR\_PKG\_POWER\_LIMIT. Two power limits can be specified, corresponding to time windows of different sizes. Each power limit provides independent clamping control that would permit the processor cores to go below OS-requested state to meet the power limits. A lock mechanism allow the software agent to enforce power limit settings. Once the lock bit is set, the power limit settings are static and un-modifiable until next RESET.

The bit fields of MSR\_PKG\_POWER\_LIMIT (Figure 14-36) are:

- **Package Power Limit #1 (bits 14:0):** Sets the average power usage limit of the package domain corresponding to time window # 1. The unit of this field is specified by the “Power Units” field of MSR\_RAPL\_POWER\_UNIT.
- **Enable Power Limit #1 (bit 15):** 0 = disabled; 1 = enabled.
- **Package Clamping Limitation #1 (bit 16):** Allow going below OS-requested P/T state setting during time window specified by bits 23:17.
- **Time Window for Power Limit #1 (bits 23:17):** Indicates the time window for power limit # 1  

$$\text{Time limit} = 2^Y * (1.0 + Z/4.0) * \text{Time\_Unit}$$

Here “Y” is the unsigned integer value represented. by bits 21:17, “Z” is an unsigned integer represented by bits 23:22. “Time\_Unit” is specified by the “Time Units” field of MSR\_RAPL\_POWER\_UNIT.



- **Package Power Limit #2(bits 46:32):** Sets the average power usage limit of the package domain corresponding to time window # 2. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Enable Power Limit #2(bit 47):** 0 = disabled; 1 = enabled.
- **Package Clamping Limitation #2 (bit 48):** Allow going below OS-requested P/T state setting during time window specified by bits 23:17.
- **Time Window for Power Limit #2 (bits 55:49):** Indicates the time window for power limit #2  

$$\text{Time limit} = 2^Y * (1.0 + Z/4.0) * \text{Time\_Unit}$$

Here "Y" is the unsigned integer value represented. by bits 53:49, "Z" is an unsigned integer represented by bits 55:54. "Time\_Unit" is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT. This field may have a hard-coded value in hardware and ignores values written by software.
- **Lock (bit 63):** If set, all write attempts to this MSR are ignored until next RESET.

MSR\_PKG\_ENERGY\_STATUS is a read-only MSR. It reports the actual energy use for the package domain. This MSR is updated every ~1msec. It has a wraparound time of around 60 secs when power consumption is high, and may be longer otherwise.




Figure 14-37. MSR\_PKG\_ENERGY\_STATUS MSR

- **Total Energy Consumed (bits 31:0):** The unsigned integer value represents the total amount of energy consumed since that last time this register is cleared. The unit of this field is specified by the "Energy Status Units" field of MSR\_RAPL\_POWER\_UNIT.

MSR\_PKG\_POWER\_INFO is a read-only MSR. It reports the package power range information for RAPL usage. This MSR provides maximum/minimum values (derived from electrical specification), thermal specification power of the package domain. It also provides the largest possible time window for software to program the RAPL interface.





Figure 14-38. MSR\_PKG\_POWER\_INFO Register

- **Thermal Spec Power (bits 14:0):** The unsigned integer value is the equivalent of thermal specification power of the package domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Minimum Power (bits 30:16):** The unsigned integer value is the equivalent of minimum power derived from electrical spec of the package domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Maximum Power (bits 46:32):** The unsigned integer value is the equivalent of maximum power derived from the electrical spec of the package domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.

- **Maximum Time Window (bits 53:48):** The unsigned integer value is the equivalent of largest acceptable value to program the time window of MSR\_PKG\_POWER\_LIMIT. The unit of this field is specified by the “Time Units” field of MSR\_RAPL\_POWER\_UNIT.

MSR\_PKG\_PERF\_STATUS is a read-only MSR. It reports the total time for which the package was throttled due to the RAPL power limits. Throttling in this context is defined as going below the OS-requested P-state or T-state. It has a wrap-around time of many hours. The availability of this MSR is platform specific (see Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*).



**Figure 14-39. MSR\_PKG\_PERF\_STATUS MSR**


- **Accumulated Package Throttled Time (bits 31:0):** The unsigned integer value represents the cumulative time (since the last time this register is cleared) that the package has throttled. The unit of this field is specified by the “Time Units” field of MSR\_RAPL\_POWER\_UNIT.

#### 14.9.4 PP0/PP1 RAPL Domains

The MSR interfaces defined for the PP0 and PP1 domains are identical in layout. Generally, PP0 refers to the processor cores. The availability of PP1 RAPL domain interface is platform-specific. For a client platform, the PP1 domain refers to the power plane of a specific device in the uncore. For server platforms, the PP1 domain is not supported, but its PP0 domain supports the MSR\_PP0\_PERF\_STATUS interface.

- MSR\_PP0\_POWER\_LIMIT/MSR\_PP1\_POWER\_LIMIT allow software to set power limits for the respective power plane domain.
- MSR\_PP0\_ENERGY\_STATUS/MSR\_PP1\_ENERGY\_STATUS report actual energy usage on a power plane.
- MSR\_PP0\_POLICY/MSR\_PP1\_POLICY allow software to adjust balance for respective power plane.

MSR\_PP0\_PERF\_STATUS can report the performance impact of power limiting, but it is not available in client platforms.



**Figure 14-40. MSR\_PP0\_POWER\_LIMIT/MSR\_PP1\_POWER\_LIMIT Register**

MSR\_PP0\_POWER\_LIMIT/MSR\_PP1\_POWER\_LIMIT allow a software agent to define power limitation for the respective power plane domain. A lock mechanism in each power plane domain allows the software agent to enforce power limit settings independently. Once a lock bit is set, the power limit settings in that power plane are static and un-modifiable until next RESET.

The bit fields of MSR\_PP0\_POWER\_LIMIT/MSR\_PP1\_POWER\_LIMIT (Figure 14-40) are:

- **Power Limit (bits 14:0):** Sets the average power usage limit of the respective power plane domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Enable Power Limit (bit 15):** 0 = disabled; 1 = enabled.
- **Clamping Limitation (bit 16):** Allow going below OS-requested P/T state setting during time window specified by bits 23:17.
- **Time Window for Power Limit (bits 23:17):** Indicates the length of time window over which the power limit #1 will be used by the processor. The numeric value encoded by bits 23:17 is represented by the product of  $2^Y * F$ ; where F is a single-digit decimal floating-point value between 1.0 and 1.3 with the fraction digit represented by bits 23:22, Y is an unsigned integer represented by bits 21:17. The unit of this field is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Lock (bit 31):** If set, all write attempts to the MSR and corresponding policy MSR\_PP0\_POLICY/MSR\_PP1\_POLICY are ignored until next RESET.

MSR\_PP0\_ENERGY\_STATUS/MSR\_PP1\_ENERGY\_STATUS are read-only MSRs. They report the actual energy use for the respective power plane domains. These MSRs are updated every ~1msec.




Figure 14-41. MSR\_PP0\_ENERGY\_STATUS/MSR\_PP1\_ENERGY\_STATUS MSR

- **Total Energy Consumed (bits 31:0):** The unsigned integer value represents the total amount of energy consumed since the last time this register was cleared. The unit of this field is specified by the "Energy Status Units" field of MSR\_RAPL\_POWER\_UNIT.

MSR\_PP0\_POLICY/MSR\_PP1\_POLICY provide balance power policy control for each power plane by providing inputs to the power budgeting management algorithm. On platforms that support PP0 (IA cores) and PP1 (uncore graphic device), the default values give priority to the non-IA power plane. These MSRs enable the PCU to balance power consumption between the IA cores and uncore graphic device.




Figure 14-42. MSR\_PP0\_POLICY/MSR\_PP1\_POLICY Register

- **Priority Level (bits 4:0):** Priority level input to the PCU for respective power plane. PP0 covers the IA processor cores, PP1 covers the uncore graphic device. The value 31 is considered highest priority.

MSR\_PP0\_PERF\_STATUS is a read-only MSR. It reports the total time for which the PP0 domain was throttled due to the power limits. This MSR is supported only in server platform. Throttling in this context is defined as going below the OS-requested P-state or T-state.




Figure 14-43. MSR\_PPO\_PERF\_STATUS MSR

- **Accumulated PPO Throttled Time (bits 31:0):** The unsigned integer value represents the cumulative time (since the last time this register is cleared) that the PPO domain has throttled. The unit of this field is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT.

### 14.9.5 DRAM RAPL Domain

The MSR interfaces defined for the DRAM domains are supported only in the server platform. The MSR interfaces are:

- MSR\_DRAM\_POWER\_LIMIT allows software to set power limits for the DRAM domain and measurement attributes associated with each limit.
- MSR\_DRAM\_ENERGY\_STATUS reports measured actual energy usage.
- MSR\_DRAM\_POWER\_INFO reports the DRAM domain power range information for RAPL usage.
- MSR\_DRAM\_PERF\_STATUS can report the performance impact of power limiting.




Figure 14-44. MSR\_DRAM\_POWER\_LIMIT Register

MSR\_DRAM\_POWER\_LIMIT allows a software agent to define power limitation for the DRAM domain. Power limitation is defined in terms of average power usage (Watts) over a time window specified in MSR\_DRAM\_POWER\_LIMIT. A power limit can be specified along with a time window. A lock mechanism allow the software agent to enforce power limit settings. Once the lock bit is set, the power limit settings are static and unmodifiable until next RESET.

The bit fields of MSR\_DRAM\_POWER\_LIMIT (Figure 14-44) are:

- **DRAM Power Limit #1 (bits 14:0):** Sets the average power usage limit of the DRAM domain corresponding to time window # 1. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Enable Power Limit #1 (bit 15):** 0 = disabled; 1 = enabled.
- **Time Window for Power Limit (bits 23:17):** Indicates the length of time window over which the power limit will be used by the processor. The numeric value encoded by bits 23:17 is represented by the product of  $2^Y * F$ ; where F is a single-digit decimal floating-point value between 1.0 and 1.3 with the fraction digit represented by bits 23:22, Y is an unsigned integer represented by bits 21:17. The unit of this field is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Lock (bit 31):** If set, all write attempts to this MSR are ignored until next RESET.

MSR\_DRAM\_ENERGY\_STATUS is a read-only MSR. It reports the actual energy use for the DRAM domain. This MSR is updated every ~1msec.




Figure 14-45. MSR\_DRAM\_ENERGY\_STATUS MSR

- **Total Energy Consumed (bits 31:0):** The unsigned integer value represents the total amount of energy consumed since that last time this register is cleared. The unit of this field is specified by the "Energy Status Units" field of MSR\_RAPL\_POWER\_UNIT.

MSR\_DRAM\_POWER\_INFO is a read-only MSR. It reports the DRAM power range information for RAPL usage. This MSR provides maximum/minimum values (derived from electrical specification), thermal specification power of the DRAM domain. It also provides the largest possible time window for software to program the RAPL interface.




Figure 14-46. MSR\_DRAM\_POWER\_INFO Register

- **Thermal Spec Power (bits 14:0):** The unsigned integer value is the equivalent of thermal specification power of the DRAM domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Minimum Power (bits 30:16):** The unsigned integer value is the equivalent of minimum power derived from electrical spec of the DRAM domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Maximum Power (bits 46:32):** The unsigned integer value is the equivalent of maximum power derived from the electrical spec of the DRAM domain. The unit of this field is specified by the "Power Units" field of MSR\_RAPL\_POWER\_UNIT.
- **Maximum Time Window (bits 53:48):** The unsigned integer value is the equivalent of largest acceptable value to program the time window of MSR\_DRAM\_POWER\_LIMIT. The unit of this field is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT.

MSR\_DRAM\_PERF\_STATUS is a read-only MSR. It reports the total time for which the package was throttled due to the RAPL power limits. Throttling in this context is defined as going below the OS-requested P-state or T-state. It has a wrap-around time of many hours. The availability of this MSR is platform specific (see Chapter 2, "Model-Specific Registers (MSRs)" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*).




Figure 14-47. MSR\_DRAM\_PERF\_STATUS MSR

- **Accumulated Package Throttled Time (bits 31:0):** The unsigned integer value represents the cumulative time (since the last time this register is cleared) that the DRAM domain has throttled. The unit of this field is specified by the "Time Units" field of MSR\_RAPL\_POWER\_UNIT.



## 11. Updates to Chapter 18, Volume 3B

Change bars show changes to Chapter 18 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2*.

-----  
Changes to this chapter: Updates to add support for 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series to existing sections. Figure and typo corrections.



Intel 64 and IA-32 architectures provide facilities for monitoring performance via a PMU (Performance Monitoring Unit).

## 18.1 PERFORMANCE MONITORING OVERVIEW

Performance monitoring was introduced in the Pentium processor with a set of model-specific performance-monitoring counter MSRs. These counters permit selection of processor performance parameters to be monitored and measured. The information obtained from these counters can be used for tuning system and compiler performance.

In Intel P6 family of processors, the performance monitoring mechanism was enhanced to permit a wider selection of events to be monitored and to allow greater control events to be monitored. Next, Intel processors based on Intel NetBurst microarchitecture introduced a distributed style of performance monitoring mechanism and performance events.

The performance monitoring mechanisms and performance events defined for the Pentium, P6 family, and Intel processors based on Intel NetBurst microarchitecture are not architectural. They are all model specific (not compatible among processor families). Intel Core Solo and Intel Core Duo processors support a set of architectural performance events and a set of non-architectural performance events. Newer Intel processor generations support enhanced architectural performance events and non-architectural performance events.

Starting with Intel Core Solo and Intel Core Duo processors, there are two classes of performance monitoring capabilities. The first class supports events for monitoring performance using counting or interrupt-based event sampling usage. These events are non-architectural and vary from one processor model to another. They are similar to those available in Pentium M processors. These non-architectural performance monitoring events are specific to the microarchitecture and may change with enhancements. They are discussed in Section 18.6.3, "Performance Monitoring (Processors Based on Intel NetBurst® Microarchitecture)." Non-architectural events for a given microarchitecture cannot be enumerated using CPUID; and they are listed in Chapter 19, "Performance Monitoring Events."

The second class of performance monitoring capabilities is referred to as architectural performance monitoring. This class supports the same counting and interrupt-based event sampling usages, with a smaller set of available events. The visible behavior of architectural performance events is consistent across processor implementations. Availability of architectural performance monitoring capabilities is enumerated using the CPUID.0AH. These events are discussed in Section 18.2.

See also:

- Section 18.2, "Architectural Performance Monitoring"
- Section 18.3, "Performance Monitoring (Intel® Core™ Processors and Intel® Xeon® Processors)"
  - Section 18.3.1, "Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Nehalem"
  - Section 18.3.2, "Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Westmere"
  - Section 18.3.3, "Intel® Xeon® Processor E7 Family Performance Monitoring Facility"
  - Section 18.3.4, "Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Sandy Bridge"
  - Section 18.3.5, "3rd Generation Intel® Core™ Processor Performance Monitoring Facility"
  - Section 18.3.6, "4th Generation Intel® Core™ Processor Performance Monitoring Facility"
  - Section 18.3.7, "5th Generation Intel® Core™ Processor and Intel® Core™ M Processor Performance Monitoring Facility"

- Section 18.3.8, “6th Generation, 7th Generation and 8th Generation Intel® Core™ Processor Performance Monitoring Facility”
- Section 18.4, “Performance monitoring (Intel® Xeon™ Phi Processors)”
  - Section 18.4.1, “Intel® Xeon Phi™ Processor 7200/5200/3200 Performance Monitoring”
- Section 18.5, “Performance Monitoring (Intel® Atom™ Processors)”
  - Section 18.5.1, “Performance Monitoring (45 nm and 32 nm Intel® Atom™ Processors)”
  - Section 18.5.2, “Performance Monitoring for Silvermont Microarchitecture”
  - Section 18.5.3, “Performance Monitoring for Goldmont Microarchitecture”
  - Section 18.5.4, “Performance Monitoring for Goldmont Plus Microarchitecture”
- Section 18.6, “Performance Monitoring (Legacy Intel Processors)”
  - Section 18.6.1, “Performance Monitoring (Intel® Core™ Solo and Intel® Core™ Duo Processors)”
  - Section 18.6.2, “Performance Monitoring (Processors Based on Intel® Core™ Microarchitecture)”
  - Section 18.6.3, “Performance Monitoring (Processors Based on Intel NetBurst® Microarchitecture)”
  - Section 18.6.4, “Performance Monitoring and Intel Hyper-Threading Technology in Processors Based on Intel NetBurst® Microarchitecture”
    - Section 18.6.4.5, “Counting Clocks on systems with Intel Hyper-Threading Technology in Processors Based on Intel NetBurst® Microarchitecture”
  - Section 18.6.5, “Performance Monitoring and Dual-Core Technology”
  - Section 18.6.6, “Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache”
  - Section 18.6.7, “Performance Monitoring on L3 and Caching Bus Controller Sub-Systems”
  - Section 18.6.8, “Performance Monitoring (P6 Family Processor)”
  - Section 18.6.9, “Performance Monitoring (Pentium Processors)”
- Section 18.7, “Counting Clocks”
- Section 18.8, “IA32\_PERF\_CAPABILITIES MSR Enumeration”

## 18.2 ARCHITECTURAL PERFORMANCE MONITORING

Performance monitoring events are architectural when they behave consistently across microarchitectures. Intel Core Solo and Intel Core Duo processors introduced architectural performance monitoring. The feature provides a mechanism for software to enumerate performance events and provides configuration and counting facilities for events.

Architectural performance monitoring does allow for enhancement across processor implementations. The CPUID.0AH leaf provides version ID for each enhancement. Intel Core Solo and Intel Core Duo processors support base level functionality identified by version ID of 1. Processors based on Intel Core microarchitecture support, at a minimum, the base level functionality of architectural performance monitoring. Intel Core 2 Duo processor T 7700 and newer processors based on Intel Core microarchitecture support both the base level functionality and enhanced architectural performance monitoring identified by version ID of 2.

45 nm and 32 nm Intel Atom processors and Intel Atom processors based on the Silvermont microarchitecture support the functionality provided by versionID 1, 2, and 3; CPUID.0AH:EAX[7:0] reports versionID = 3 to indicate the aggregate of architectural performance monitoring capabilities. Intel Atom processors based on the Airmont microarchitecture support the same performance monitoring capabilities as those based on the Silvermont microarchitecture.

Intel Core processors and related Intel Xeon processor families based on the Nehalem through Broadwell microarchitectures support version ID 1, 2, and 3. Intel processors based on the Skylake, Kaby Lake and Coffee Lake microarchitectures support versionID 4.

Next generation Intel Atom processors are based on the Goldmont microarchitecture. Intel processors based on the Goldmont microarchitecture support versionID 4.

## 18.2.1 Architectural Performance Monitoring Version 1

Configuring an architectural performance monitoring event involves programming performance event select registers. There are a finite number of performance event select MSRs (IA32\_PERFEVTSELx MSRs). The result of a performance monitoring event is reported in a performance monitoring counter (IA32\_PMCx MSR). Performance monitoring counters are paired with performance monitoring select registers.

Performance monitoring select registers and counters are architectural in the following respects:

- Bit field layout of IA32\_PERFEVTSELx is consistent across microarchitectures.
- Addresses of IA32\_PERFEVTSELx MSRs remain the same across microarchitectures.
- Addresses of IA32\_PMC MSRs remain the same across microarchitectures.
- Each logical processor has its own set of IA32\_PERFEVTSELx and IA32\_PMCx MSRs. Configuration facilities and counters are not shared between logical processors sharing a processor core.

Architectural performance monitoring provides a CPUID mechanism for enumerating the following information:

- Number of performance monitoring counters available in a logical processor (each IA32\_PERFEVTSELx MSR is paired to the corresponding IA32\_PMCx MSR)
- Number of bits supported in each IA32\_PMCx
- Number of architectural performance monitoring events supported in a logical processor

Software can use CPUID to discover architectural performance monitoring availability (CPUID.0AH). The architectural performance monitoring leaf provides an identifier corresponding to the version number of architectural performance monitoring available in the processor.

The version identifier is retrieved by querying CPUID.0AH:EAX[bits 7:0] (see Chapter 3, "Instruction Set Reference, A-L," in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A*). If the version identifier is greater than zero, architectural performance monitoring capability is supported. Software queries the CPUID.0AH for the version identifier first; it then analyzes the value returned in CPUID.0AH.EAX, CPUID.0AH.EBX to determine the facilities available.

In the initial implementation of architectural performance monitoring; software can determine how many IA32\_PERFEVTSELx/ IA32\_PMCx MSR pairs are supported per core, the bit-width of PMC, and the number of architectural performance monitoring events available.

### 18.2.1.1 Architectural Performance Monitoring Version 1 Facilities

Architectural performance monitoring facilities include a set of performance monitoring counters and performance event select registers. These MSRs have the following properties:

- IA32\_PMCx MSRs start at address 0C1H and occupy a contiguous block of MSR address space; the number of MSRs per logical processor is reported using CPUID.0AH:EAX[15:8].
- IA32\_PERFEVTSELx MSRs start at address 186H and occupy a contiguous block of MSR address space. Each performance event select register is paired with a corresponding performance counter in the 0C1H address block.
- The bit width of an IA32\_PMCx MSR is reported using the CPUID.0AH:EAX[23:16]. This the number of valid bits for read operation. On write operations, the lower-order 32 bits of the MSR may be written with any value, and the high-order bits are sign-extended from the value of bit 31.
- Bit field layout of IA32\_PERFEVTSELx MSRs is defined architecturally.

See Figure 18-1 for the bit field layout of IA32\_PERFEVTSELx MSRs. The bit fields are:

- **Event select field (bits 0 through 7)** — Selects the event logic unit used to detect microarchitectural conditions (see Table 18-1, for a list of architectural events and their 8-bit codes). The set of values for this field is defined architecturally; each value corresponds to an event logic unit for use with an architectural performance event. The number of architectural events is queried using CPUID.0AH:EAX. A processor may support only a subset of pre-defined values.

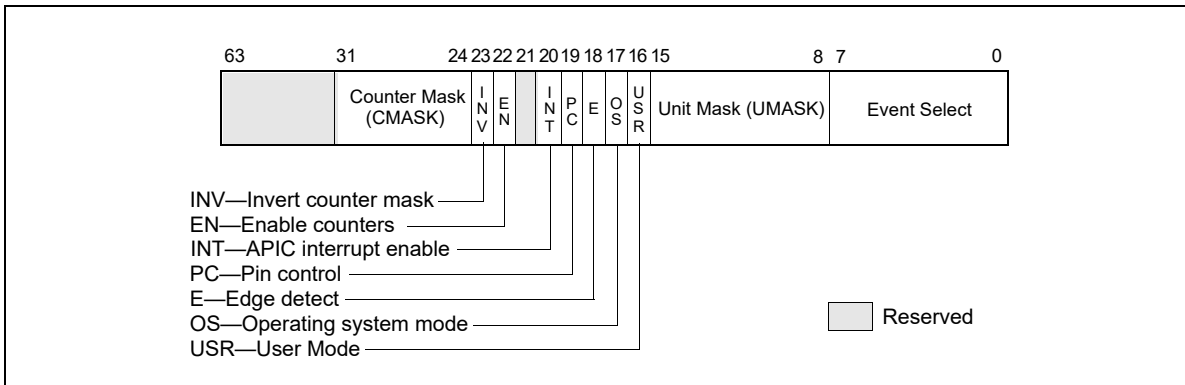


Figure 18-1. Layout of IA32\_PERFEVTSELx MSRs

- **Unit mask (UMASK) field (bits 8 through 15)** — These bits qualify the condition that the selected event logic unit detects. Valid UMASK values for each event logic unit are specific to the unit. For each architectural performance event, its corresponding UMASK value defines a specific microarchitectural condition. A pre-defined microarchitectural condition associated with an architectural event may not be applicable to a given processor. The processor then reports only a subset of pre-defined architectural events. Pre-defined architectural events are listed in Table 18-1; support for pre-defined architectural events is enumerated using CPUID.0AH:EBX. Architectural performance events available in the initial implementation are listed in Table 19-1.
- **USR (user mode) flag (bit 16)** — Specifies that the selected microarchitectural condition is counted when the logical processor is operating at privilege levels 1, 2 or 3. This flag can be used with the OS flag.
- **OS (operating system mode) flag (bit 17)** — Specifies that the selected microarchitectural condition is counted when the logical processor is operating at privilege level 0. This flag can be used with the USR flag.
- **E (edge detect) flag (bit 18)** — Enables (when set) edge detection of the selected microarchitectural condition. The logical processor counts the number of deasserted to asserted transitions for any condition that can be expressed by the other fields. The mechanism does not permit back-to-back assertions to be distinguished. This mechanism allows software to measure not only the fraction of time spent in a particular state, but also the average length of time spent in such a state (for example, the time spent waiting for an interrupt to be serviced).
- **PC (pin control) flag (bit 19)** — When set, the logical processor toggles the PM*i* pins and increments the counter when performance-monitoring events occur; when clear, the processor toggles the PM*i* pins when the counter overflows. The toggling of a pin is defined as assertion of the pin for a single bus clock followed by deassertion.
- **INT (APIC interrupt enable) flag (bit 20)** — When set, the logical processor generates an exception through its local APIC on counter overflow.
- **EN (Enable Counters) Flag (bit 22)** — When set, performance counting is enabled in the corresponding performance-monitoring counter; when clear, the corresponding counter is disabled. The event logic unit for a UMASK must be disabled by setting IA32\_PERFEVTSELx[bit 22] = 0, before writing to IA32\_PMCx.
- **INV (invert) flag (bit 23)** — When set, inverts the counter-mask (CMASK) comparison, so that both greater than or equal to and less than comparisons can be made (0: greater than or equal; 1: less than). Note if counter-mask is programmed to zero, INV flag is ignored.

- **Counter mask (CMASK) field (bits 24 through 31)** — When this field is not zero, a logical processor compares this mask to the events count of the detected microarchitectural condition during a single cycle. If the event count is greater than or equal to this mask, the counter is incremented by one. Otherwise the counter is not incremented.

This mask is intended for software to characterize microarchitectural conditions that can count multiple occurrences per cycle (for example, two or more instructions retired per clock; or bus queue occupations). If the counter-mask field is 0, then the counter is incremented each cycle by the event count associated with multiple occurrences.

### 18.2.1.2 Pre-defined Architectural Performance Events

Table 18-1 lists architecturally defined events.

**Table 18-1. UMask and Event Select Encodings for Pre-Defined Architectural Performance Events**

| Bit Position<br>CPUID.AH.EBX | Event Name                 | UMask | Event Select |
|------------------------------|----------------------------|-------|--------------|
| 0                            | UnHalted Core Cycles       | 00H   | 3CH          |
| 1                            | Instruction Retired        | 00H   | C0H          |
| 2                            | UnHalted Reference Cycles  | 01H   | 3CH          |
| 3                            | LLC Reference              | 4FH   | 2EH          |
| 4                            | LLC Misses                 | 41H   | 2EH          |
| 5                            | Branch Instruction Retired | 00H   | C4H          |
| 6                            | Branch Misses Retired      | 00H   | C5H          |

A processor that supports architectural performance monitoring may not support all the predefined architectural performance events (Table 18-1). The non-zero bits in CPUID.0AH:EBX indicate the events that are not available.

The behavior of each architectural performance event is expected to be consistent on all processors that support that event. Minor variations between microarchitectures are noted below:

- **UnHalted Core Cycles** — Event select 3CH, Umask 00H

This event counts core clock cycles when the clock signal on a specific core is running (not halted). The counter does not advance in the following conditions:

- an ACPI C-state other than C0 for normal operation
- HLT
- STPCLK# pin asserted
- being throttled by TM1
- during the frequency switching phase of a performance state transition (see Chapter 14, “Power and Thermal Management”)

The performance counter for this event counts across performance state transitions using different core clock frequencies

- **Instructions Retired** — Event select C0H, Umask 00H

This event counts the number of instructions at retirement. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. An instruction with a REP prefix counts as one instruction (not per iteration). Faults before the retirement of the last micro-op of a multi-ops instruction are not counted.

This event does not increment under VM-exit conditions. Counters continue counting during hardware interrupts, traps, and inside interrupt handlers.

- UnHalted Reference Cycles** — Event select 3CH, Umask 01H  
 This event counts reference clock cycles at a fixed frequency while the clock signal on the core is running. The event counts at a fixed frequency, irrespective of core frequency changes due to performance state transitions. Processors may implement this behavior differently. Current implementations use the core crystal clock, TSC or the bus clock. Because the rate may differ between implementations, software should calibrate it to a time source with known frequency.
- Last Level Cache References** — Event select 2EH, Umask 4FH  
 This event counts requests originating from the core that reference a cache line in the last level on-die cache. The event count includes speculation and cache line fills due to the first-level cache hardware prefetcher, but may exclude cache line fills due to other hardware-prefetchers.  
 Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended.
- Last Level Cache Misses** — Event select 2EH, Umask 41H  
 This event counts each cache miss condition for references to the last level on-die cache. The event count may include speculation and cache line fills due to the first-level cache hardware prefetcher, but may exclude cache line fills due to other hardware-prefetchers.  
 Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended.
- Branch Instructions Retired** — Event select C4H, Umask 00H  
 This event counts branch instructions at retirement. It counts the retirement of the last micro-op of a branch instruction.
- All Branch Mispredict Retired** — Event select C5H, Umask 00H  
 This event counts mispredicted branch instructions at retirement. It counts the retirement of the last micro-op of a branch instruction in the architectural path of execution and experienced misprediction in the branch prediction hardware.  
 Branch prediction hardware is implementation-specific across microarchitectures; value comparison to estimate performance differences is not recommended.

#### NOTE

Programming decisions or software precisions on functionality should not be based on the event values or dependent on the existence of performance monitoring events.

## 18.2.2 Architectural Performance Monitoring Version 2

The enhanced features provided by architectural performance monitoring version 2 include the following:

- Fixed-function performance counter register and associated control register** — Three of the architectural performance events are counted using three fixed-function MSRs (IA32\_FIXED\_CTR0 through IA32\_FIXED\_CTR2). Each of the fixed-function PMC can count only one architectural performance event.  
 Configuring the fixed-function PMCs is done by writing to bit fields in the MSR (IA32\_FIXED\_CTR\_CTRL) located at address 38DH. Unlike configuring performance events for general-purpose PMCs (IA32\_PMCx) via UMASK field in (IA32\_PERFECTSELx), configuring, programming IA32\_FIXED\_CTR\_CTRL for fixed-function PMCs do not require any UMASK.

- Simplified event programming** — Most frequent operation in programming performance events are enabling/disabling event counting and checking the status of counter overflows. Architectural performance event version 2 provides three architectural MSR:
  - IA32\_PERF\_GLOBAL\_CTRL allows software to enable/disable event counting of all or any combination of fixed-function PMCs (IA32\_FIXED\_CTRx) or any general-purpose PMCs via a single WRMSR.
  - IA32\_PERF\_GLOBAL\_STATUS allows software to query counter overflow conditions on any combination of fixed-function PMCs or general-purpose PMCs via a single RDMSR.
  - IA32\_PERF\_GLOBAL\_OVF\_CTRL allows software to clear counter overflow conditions on any combination of fixed-function PMCs or general-purpose PMCs via a single WRMSR.
- PMI Overhead Mitigation** — Architectural performance monitoring version 2 introduces two bit field interface in IA32\_DEBUGCTL for PMI service routine to accumulate performance monitoring data and LBR records with reduced perturbation from servicing the PMI. The two bit fields are:
  - IA32\_DEBUGCTL.Freeze\_LBR\_On\_PMI(bit 11). In architectural performance monitoring version 2, only the legacy semantic behavior is supported. See Section 17.4.7 for details of the legacy Freeze LBRs on PMI control.
  - IA32\_DEBUGCTL.Freeze\_PerfMon\_On\_PMI(bit 12). In architectural performance monitoring version 2, only the legacy semantic behavior is supported. See Section 17.4.7 for details of the legacy Freeze LBRs on PMI control.

The facilities provided by architectural performance monitoring version 2 can be queried from CPUID leaf 0AH by examining the content of register EDX:

- Bits 0 through 4 of CPUID.0AH.EDX indicates the number of fixed-function performance counters available per core,
- Bits 5 through 12 of CPUID.0AH.EDX indicates the bit-width of fixed-function performance counters. Bits beyond the width of the fixed-function counter are reserved and must be written as zeros.

NOTE

Early generation of processors based on Intel Core microarchitecture may report in CPUID.0AH:EDX of support for version 2 but indicating incorrect information of version 2 facilities.

The IA32\_FIXED\_CTR\_CTRL MSR include multiple sets of 4-bit field, each 4 bit field controls the operation of a fixed-function performance counter. Figure 18-2 shows the layout of 4-bit controls for each fixed-function PMC. Two sub-fields are currently defined within each control. The definitions of the bit fields are:

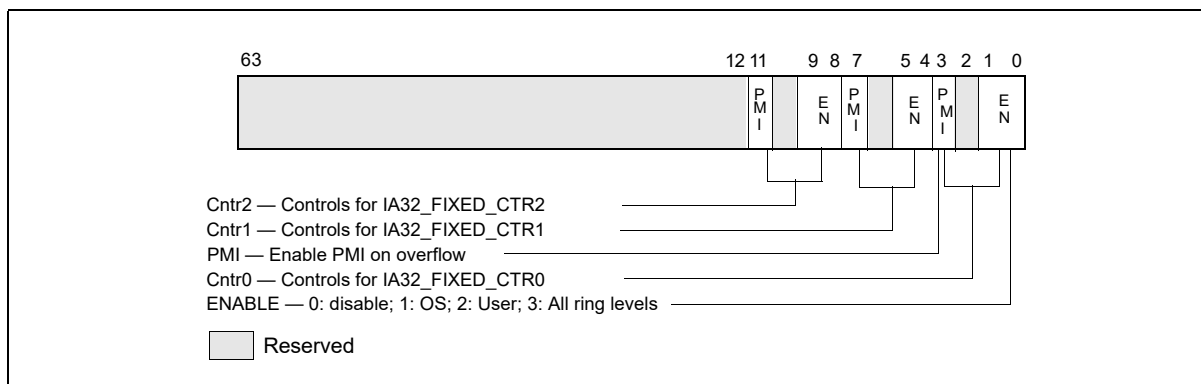



Figure 18-2. Layout of IA32\_FIXED\_CTR\_CTRL MSR

- **Enable field (lowest 2 bits within each 4-bit control)** — When bit 0 is set, performance counting is enabled in the corresponding fixed-function performance counter to increment while the target condition associated with the architecture performance event occurred at ring 0. When bit 1 is set, performance counting is enabled in the corresponding fixed-function performance counter to increment while the target condition associated with the architecture performance event occurred at ring greater than 0. Writing 0 to both bits stops the performance counter. Writing a value of 11B enables the counter to increment irrespective of privilege levels.
- **PMI field (the fourth bit within each 4-bit control)** — When set, the logical processor generates an exception through its local APIC on overflow condition of the respective fixed-function counter.

IA32\_PERF\_GLOBAL\_CTRL MSR provides single-bit controls to enable counting of each performance counter. Figure 18-3 shows the layout of IA32\_PERF\_GLOBAL\_CTRL. Each enable bit in IA32\_PERF\_GLOBAL\_CTRL is AND'ed with the enable bits for all privilege levels in the respective IA32\_PERFEVTSELx or IA32\_PERF\_FIXED\_CTR\_CTRL MSRs to start/stop the counting of respective counters. Counting is enabled if the AND'ed results is true; counting is disabled when the result is false.



**Figure 18-3. Layout of IA32\_PERF\_GLOBAL\_CTRL MSR**

The behavior of the fixed function performance counters supported by architectural performance version 2 is expected to be consistent on all processors that support those counters, and is defined as follows.



**Table 18-2. Association of Fixed-Function Performance Counters with Architectural Performance Events**

| Fixed-Function Performance Counter   | Address | Event Mask Mnemonic                              | Description   |
|--------------------------------------|---------|--|---|
| MSR_PERF_FIXED_CTR0//IA32_FIXED_CTR0 | 309H    | INST_RETIRED.ANY                                 | This event counts the number of instructions that retire execution. For instructions that consist of multiple uops, this event counts the retirement of the last uop of the instruction. The counter continues counting during hardware interrupts, traps, and in-side interrupt handlers.  |
| MSR_PERF_FIXED_CTR1//IA32_FIXED_CTR1 | 30AH    | CPU_CLK_UNHALTED.THREAD<br>CPU_CLK_UNHALTED.CORE | The CPU_CLK_UNHALTED.THREAD event counts the number of core cycles while the logical processor is not in a halt state.<br><br>If there is only one logical processor in a processor core, CPU_CLK_UNHALTED.CORE counts the unhalting cycles of the processor core.<br><br>The core frequency may change from time to time due to transitions associated with Enhanced Intel SpeedStep Technology or TM2. For this reason this event may have a changing ratio with regards to time.       |
| MSR_PERF_FIXED_CTR2//IA32_FIXED_CTR2 | 30BH    | CPU_CLK_UNHALTED.REF_TSC                         | This event counts the number of reference cycles at the TSC rate when the core is not in a halt state and not in a TM stop-clock state. The core enters the halt state when it is running the HLT instruction or the MWAIT instruction. This event is not affected by core frequency changes (e.g., P states) but counts at the same frequency as the time stamp counter. This event can approximate elapsed time while the core was not in a halt state and not in a TM stopclock state. |

IA32\_PERF\_GLOBAL\_STATUS MSR provides single-bit status for software to query the overflow condition of each performance counter. IA32\_PERF\_GLOBAL\_STATUS[bit 62] indicates overflow conditions of the DS area data buffer. IA32\_PERF\_GLOBAL\_STATUS[bit 63] provides a CondChgd bit to indicate changes to the state of performance monitoring hardware. Figure 18-4 shows the layout of IA32\_PERF\_GLOBAL\_STATUS. A value of 1 in bits 0, 1, 32 through 34 indicates a counter overflow condition has occurred in the associated counter.

When a performance counter is configured for PEBS, overflow condition in the counter generates a performance-monitoring interrupt signaling a PEBS event. On a PEBS event, the processor stores data records into the buffer area (see Section 18.15.5), clears the counter overflow status, and sets the "OvfBuffer" bit in IA32\_PERF\_GLOBAL\_STATUS.




Figure 18-4. Layout of IA32\_PERF\_GLOBAL\_STATUS MSR

IA32\_PERF\_GLOBAL\_OVF\_CTL MSR allows software to clear overflow indicator(s) of any general-purpose or fixed-function counters via a single WRMSR. Software should clear overflow indications when

- Setting up new values in the event select and/or UMASK field for counting or interrupt-based event sampling.
- Reloading counter values to continue collecting next sample.
- Disabling event counting or interrupt-based event sampling.

The layout of IA32\_PERF\_GLOBAL\_OVF\_CTL is shown in Figure 18-5.




Figure 18-5. Layout of IA32\_PERF\_GLOBAL\_OVF\_CTRL MSR

### 18.2.3 Architectural Performance Monitoring Version 3

Processors supporting architectural performance monitoring version 3 also supports version 1 and 2, as well as capability enumerated by CPUID leaf 0AH. Specifically, version 3 provides the following enhancement in performance monitoring facilities if a processor core comprising of more than one logical processor, i.e. a processor core supporting Intel Hyper-Threading Technology or simultaneous multi-threading capability:

- AnyThread counting for processor core supporting two or more logical processors. The interface that supports AnyThread counting include:
  - Each IA32\_PERFEVTSELx MSR (starting at MSR address 186H) support the bit field layout defined in Figure 18-6.

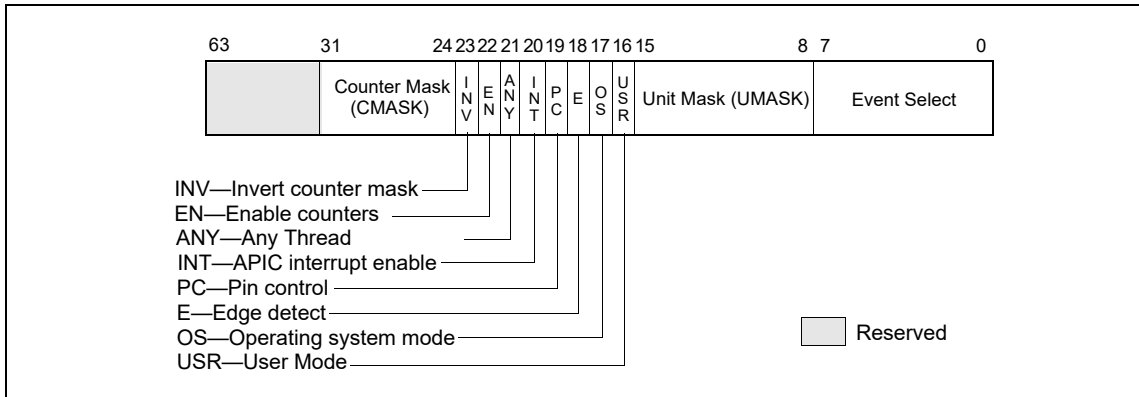


Figure 18-6. Layout of IA32\_PERFEVTSELx MSRs Supporting Architectural Performance Monitoring Version 3

Bit 21 (AnyThread) of IA32\_PERFEVTSELx is supported in architectural performance monitoring version 3 for processor core comprising of two or more logical processors. When set to 1, it enables counting the associated event conditions (including matching the thread’s CPL with the OS/USR setting of IA32\_PERFEVTSELx) occurring across all logical processors sharing a processor core. When bit 21 is 0, the counter only increments the associated event conditions (including matching the thread’s CPL with the OS/USR setting of IA32\_PERFEVTSELx) occurring in the logical processor which programmed the IA32\_PERFEVTSELx MSR.

- Each fixed-function performance counter IA32\_FIXED\_CTRx (starting at MSR address 309H) is configured by a 4-bit control block in the IA32\_PERF\_FIXED\_CTR\_CTRL MSR. The control block also allow thread-specificity configuration using an AnyThread bit. The layout of IA32\_PERF\_FIXED\_CTR\_CTRL MSR is shown.

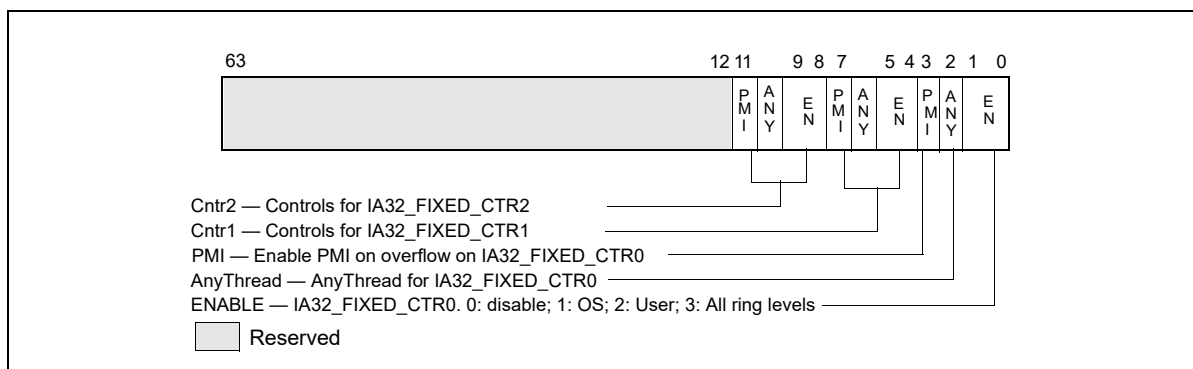
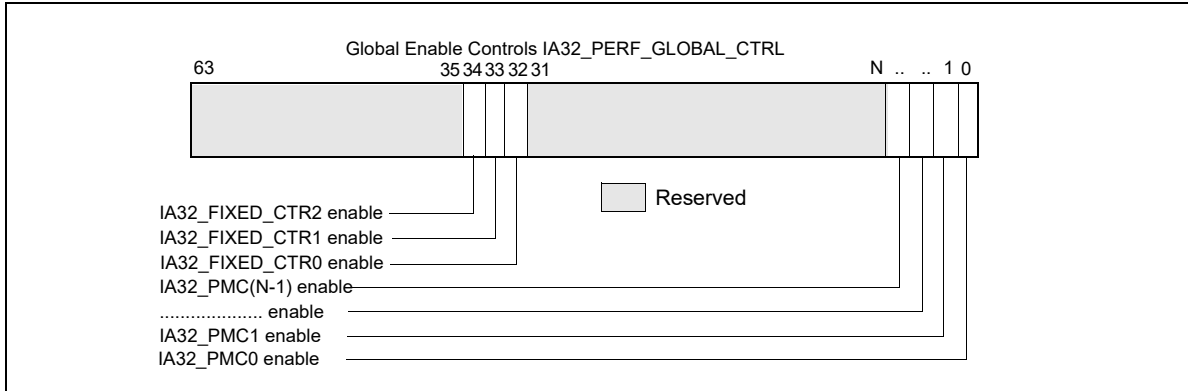


Figure 18-7. IA32\_PERF\_FIXED\_CTR\_CTRL MSR Supporting Architectural Performance Monitoring Version 3


Each control block for a fixed-function performance counter provides a **AnyThread** (bit position 2 + 4\*N, N= 0, 1, etc.) bit. When set to 1, it enables counting the associated event conditions (including matching the thread’s CPL with the ENABLE setting of the corresponding control block of IA32\_PERF\_FIXED\_CTR\_CTRL) occurring across all logical processors sharing a processor core. When an **AnyThread** bit is 0 in IA32\_PERF\_FIXED\_CTR\_CTRL, the corresponding fixed counter only increments the associated event conditions occurring in the logical processor which programmed the IA32\_PERF\_FIXED\_CTR\_CTRL MSR.

- The IA32\_PERF\_GLOBAL\_CTRL, IA32\_PERF\_GLOBAL\_STATUS, IA32\_PERF\_GLOBAL\_OVF\_CTRL MSRs provide single-bit controls/status for each general-purpose and fixed-function performance counter. Figure 18-8 and Figure 18-9 show the layout of these MSRs for N general-purpose performance counters (where N is reported by CPUID.0AH:EAX[15:8]) and three fixed-function counters.

**Note:** The number of general-purpose performance monitoring counters (i.e. N in Figure 18-9) can vary across processor generations within a processor family, across processor families, or could be different depending on the configuration chosen at boot time in the BIOS regarding Intel Hyper Threading Technology, (e.g. N=2 for 45 nm Intel Atom processors; N =4 for processors based on the Nehalem microarchitecture; for processors based on the Sandy Bridge microarchitecture, N = 4 if Intel Hyper Threading Technology is active and N=8 if not active).



**Figure 18-8. Layout of Global Performance Monitoring Control MSR**



**Figure 18-9. Global Performance Monitoring Overflow Status and Control MSRs**

### 18.2.3.1 AnyThread Counting and Software Evolution

The motivation for characterizing software workload over multiple software threads running on multiple logical processors of the same processor core originates from a time earlier than the introduction of the AnyThread interface in IA32\_PERFEVTSELx and IA32\_FIXED\_CTR\_CTRL. While AnyThread counting provides some benefits in simple software environments of an earlier era, the evolution contemporary software environments introduce certain concepts and pre-requisites that AnyThread counting does not comply with.

One example is the proliferation of software environments that support multiple virtual machines (VM) under VMX (see Chapter 23, “Introduction to Virtual-Machine Extensions”) where each VM represents a domain separated from one another.

A Virtual Machine Monitor (VMM) that manages the VMs may allow individual VM to employ performance monitoring facilities to profiles the performance characteristics of a workload. The use of the Anythread interface in IA32\_PERFEVTSELx and IA32\_FIXED\_CTR\_CTRL is discouraged with software environments supporting virtualization or requiring domain separation.

Specifically, Intel recommends VMM:

- configure the MSR bitmap to cause VM-exits for WRMSR to IA32\_PERFEVTSELx and IA32\_FIXED\_CTR\_CTRL in VMX non-Root operation (see CHAPTER 24 for additional information),
- clear the AnyThread bit of IA32\_PERFEVTSELx and IA32\_FIXED\_CTR\_CTRL in the MSR-load lists for VM exits and VM entries (see CHAPTER 24, CHAPTER 26, and CHAPTER 27).

Even when operating in simpler legacy software environments which might not emphasize the pre-requisites of a virtualized software environment, the use of the AnyThread interface should be moderated and follow any event-specific guidance where explicitly noted (see relevant sections of Chapter 19, “Performance Monitoring Events”).

## 18.2.4 Architectural Performance Monitoring Version 4

Processors supporting architectural performance monitoring version 4 also supports version 1, 2, and 3, as well as capability enumerated by CPUID leaf 0AH. Version 4 introduced a streamlined PMI overhead mitigation interface that replaces the legacy semantic behavior but retains the same control interface in IA32\_DEBUGCTL.Freeze\_LBRs\_On\_PMI and Freeze\_PerfMon\_On\_PMI. Specifically version 4 provides the following enhancement:

- New indicators (LBR\_FRZ, CTR\_FRZ) in IA32\_PERF\_GLOBAL\_STATUS, see Section 18.2.4.1.
- Streamlined Freeze/PMI Overhead management interfaces to use IA32\_DEBUGCTL.Freeze\_LBRs\_On\_PMI and IA32\_DEBUGCTL.Freeze\_PerfMon\_On\_PMI: see Section 18.2.4.1. Legacy semantics of Freeze\_LBRs\_On\_PMI and Freeze\_PerfMon\_On\_PMI (applicable to version 2 and 3) are not supported with version 4 or higher.
- Fine-grain separation of control interface to manage overflow/status of IA32\_PERF\_GLOBAL\_STATUS and read-only performance counter enabling interface in IA32\_PERF\_GLOBAL\_STATUS: see Section 18.2.4.2.
- Performance monitoring resource in-use MSR to facilitate cooperative sharing protocol between perfmon-managing privilege agents.

### 18.2.4.1 Enhancement in IA32\_PERF\_GLOBAL\_STATUS

The IA32\_PERF\_GLOBAL\_STATUS MSR provides the following indicators with architectural performance monitoring version 4:

- IA32\_PERF\_GLOBAL\_STATUS.LBR\_FRZ[bit 58]: This bit is set due to the following conditions:
  - IA32\_DEBUGCTL.FREEZE\_LBR\_ON\_PMI has been set by the profiling agent, and
  - A performance counter, configured to generate PMI, has overflowed to signal a PMI. Consequently the LBR stack is frozen.

Effectively, the IA32\_PERF\_GLOBAL\_STATUS.LBR\_FRZ bit also serve as an read-only control to enable capturing data in the LBR stack. To enable capturing LBR records, the following expression must hold with architectural perfmon version 4 or higher:

- $(\text{IA32\_DEBUGCTL.LBR} \ \& \ (!\text{IA32\_PERF\_GLOBAL\_STATUS.LBR\_FRZ})) = 1$

- IA32\_PERF\_GLOBAL\_STATUS.CTR\_FRZ[bit 59]: This bit is set due to the following conditions:
  - IA32\_DEBUGCTL.FREEZE\_PERFMON\_ON\_PMI has been set by the profiling agent, and
  - A performance counter, configured to generate PMI, has overflowed to signal a PMI. Consequently, all the performance counters are frozen.

Effectively, the IA32\_PERF\_GLOBAL\_STATUS.CTR\_FRZ bit also serve as an read-only control to enable programmable performance counters and fixed counters in the core PMU. To enable counting with the performance counters, the following expression must hold with architectural perfmon version 4 or higher:

- $(IA32\_PERFEVTSELn.EN \ \& \ IA32\_PERF\_GLOBAL\_CTRL.PMCn \ \& \ (!IA32\_PERF\_GLOBAL\_STATUS.CTR\_FRZ)) = 1$  for programmable counter 'n', or
- $(IA32\_PERF\_FIXED\_CTRL.ENi \ \& \ IA32\_PERF\_GLOBAL\_CTRL.FCi \ \& \ (!IA32\_PERF\_GLOBAL\_STATUS.CTR\_FRZ)) = 1$  for fixed counter 'i'

The read-only enable interface IA32\_PERF\_GLOBAL\_STATUS.CTR\_FRZ provides a more efficient flow for a PMI handler to use IA32\_DEBUGCTL.Freeza\_Perfmon\_On\_PMI to filter out data that may distort target workload analysis, see Table 17-3. It should be noted the IA32\_PERF\_GLOBAL\_CTRL register continue to serve as the primary interface to control all performance counters of the logical processor.

For example, when the Freeze-On-PMI mode is not being used, a PMI handler would be setting IA32\_PERF\_GLOBAL\_CTRL as the very last step to commence the overall operation after configuring the individual counter registers, controls and PEBS facility. This does not only assure atomic monitoring but also avoids unnecessary complications (e.g. race conditions) when software attempts to change the core PMU configuration while some counters are kept enabled.

Additionally, IA32\_PERF\_GLOBAL\_STATUS.TraceToPAPMI[bit 55]: On processors that support Intel Processor Trace and configured to store trace output packets to physical memory using the ToPA scheme, bit 55 is set when a PMI occurred due to a ToPA entry memory buffer was completely filled.

IA32\_PERF\_GLOBAL\_STATUS also provides an indicator to distinguish interaction of performance monitoring operations with other side-band activities, which apply Intel SGX on processors that support SGX (For additional information about Intel SGX, see "Intel® Software Guard Extensions Programming Reference".):

- IA32\_PERF\_GLOBAL\_STATUS.ASCI[bit 60]: This bit is set when data accumulated in any of the configured performance counters (i.e. IA32\_PMCx or IA32\_FIXED\_CTRx) may include contributions from direct or indirect operation of Intel SGX to protect an enclave (since the last time IA32\_PERF\_GLOBAL\_STATUS.ASCI was cleared).




Figure 18-10. IA32\_PERF\_GLOBAL\_STATUS MSR and Architectural Perfmon Version 4

Note, a processor’s support for IA32\_PERF\_GLOBAL\_STATUS.TraceToPAPMI[bit 55] is enumerated as a result of CPUID enumerated capability of Intel Processor Trace and the use of the ToPA buffer scheme. Support of IA32\_PERF\_GLOBAL\_STATUS.ASCI[bit 60] is enumerated by the CPUID enumeration of Intel SGX.

### 18.2.4.2 IA32\_PERF\_GLOBAL\_STATUS\_RESET and IA32\_PERF\_GLOBAL\_STATUS\_SET MSRS

With architectural performance monitoring version 3 and lower, clearing of the set bits in IA32\_PERF\_GLOBAL\_STATUS MSR by software is done via IA32\_PERF\_GLOBAL\_OVF\_CTRL MSR. Starting with architectural performance monitoring version 4, software can manage the overflow and other indicators in IA32\_PERF\_GLOBAL\_STATUS using separate interfaces to set or clear individual bits.

The address and the architecturally-defined bits of IA32\_PERF\_GLOBAL\_OVF\_CTRL is inherited by IA32\_PERF\_GLOBAL\_STATUS\_RESET (see Figure 18-11). Further, IA32\_PERF\_GLOBAL\_STATUS\_RESET provides additional bit fields to clear the new indicators in IA32\_PERF\_GLOBAL\_STATUS described in Section 18.2.4.1.

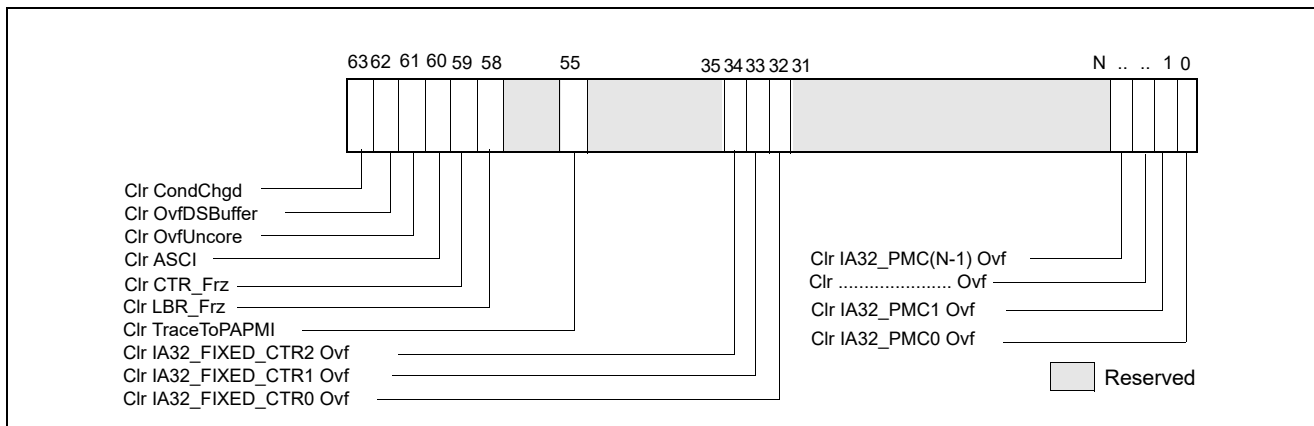


Figure 18-11. IA32\_PERF\_GLOBAL\_STATUS\_RESET MSR and Architectural Perfmon Version 4

The IA32\_PERF\_GLOBAL\_STATUS\_SET MSR is introduced with architectural performance monitoring version 4. It allows software to set individual bits in IA32\_PERF\_GLOBAL\_STATUS. The IA32\_PERF\_GLOBAL\_STATUS\_SET interface can be used by a VMM to virtualize the state of IA32\_PERF\_GLOBAL\_STATUS across VMs.

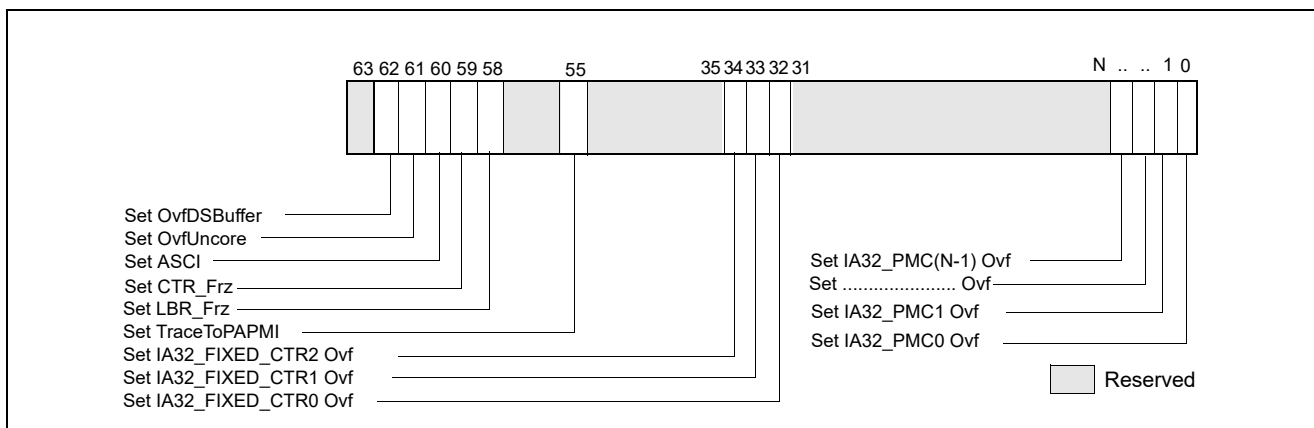


Figure 18-12. IA32\_PERF\_GLOBAL\_STATUS\_SET MSR and Architectural Perfmon Version 4

### 18.2.4.3 IA32\_PERF\_GLOBAL\_INUSE MSR

In a contemporary software environment, multiple privileged service agents may wish to employ the processor’s performance monitoring facilities. The IA32\_MISC\_ENABLE.PERFMON\_AVAILABLE[bit 7] interface could not serve

the need of multiple agent adequately. A white paper, “Performance Monitoring Unit Sharing Guideline”<sup>1</sup>, proposed a cooperative sharing protocol that is voluntary for participating software agents.

Architectural performance monitoring version 4 introduces a new MSR, IA32\_PERF\_GLOBAL\_INUSE, that simplifies the task of multiple cooperating agents to implement the sharing protocol.

The layout of IA32\_PERF\_GLOBAL\_INUSE is shown in Figure 18-13.




Figure 18-13. IA32\_PERF\_GLOBAL\_INUSE MSR and Architectural Perfmon Version 4

The IA32\_PERF\_GLOBAL\_INUSE MSR provides an “InUse” bit for each programmable performance counter and fixed counter in the processor. Additionally, it includes an indicator if the PMI mechanism has been configured by a profiling agent.

- IA32\_PERF\_GLOBAL\_INUSE.PERFEVTSEL0\_InUse[bit 0]: This bit reflects the logical state of (IA32\_PERFEVTSEL0[7:0] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.PERFEVTSEL1\_InUse[bit 1]: This bit reflects the logical state of (IA32\_PERFEVTSEL1[7:0] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.PERFEVTSEL2\_InUse[bit 2]: This bit reflects the logical state of (IA32\_PERFEVTSEL2[7:0] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.PERFEVTSELn\_InUse[bit n]: This bit reflects the logical state of (IA32\_PERFEVTSELn[7:0] != 0), n < CPUID.0AH:EAX[15:8].
- IA32\_PERF\_GLOBAL\_INUSE.FC0\_InUse[bit 32]: This bit reflects the logical state of (IA32\_FIXED\_CTR\_CTRL[1:0] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.FC1\_InUse[bit 33]: This bit reflects the logical state of (IA32\_FIXED\_CTR\_CTRL[5:4] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.FC2\_InUse[bit 34]: This bit reflects the logical state of (IA32\_FIXED\_CTR\_CTRL[9:8] != 0).
- IA32\_PERF\_GLOBAL\_INUSE.PMI\_InUse[bit 63]: This bit is set if any one of the following bit is set:
  - IA32\_PERFEVTSELn.INT[bit 20], n < CPUID.0AH:EAX[15:8].
  - IA32\_FIXED\_CTR\_CTRL.ENi\_PMI, i = 0, 1, 2.
  - Any IA32\_PEBS\_ENABLES bit which enables PEBS for a general-purpose or fixed-function performance counter.

1. Available at <http://www.intel.com/sdm>



## 18.2.5 Full-Width Writes to Performance Counter Registers

The general-purpose performance counter registers IA32\_PMCx are writable via WRMSR instruction. However, the value written into IA32\_PMCx by WRMSR is the signed extended 64-bit value of the EAX[31:0] input of WRMSR.

A processor that supports full-width writes to the general-purpose performance counters enumerated by CPUID.0AH:EAX[15:8] will set IA32\_PERF\_CAPABILITIES[13] to enumerate its full-width-write capability. See Figure 18-63.

If IA32\_PERF\_CAPABILITIES.FW\_WRITE[bit 13] = 1, each IA32\_PMCi is accompanied by a corresponding alias address starting at 4C1H for IA32\_A\_PMC0.

The bit width of the performance monitoring counters is specified in CPUID.0AH:EAX[23:16].

If IA32\_A\_PMCi is present, the 64-bit input value (EDX:EAX) of WRMSR to IA32\_A\_PMCi will cause IA32\_PMCi to be updated by:

```
COUNTERWIDTH = CPUID.0AH:EAX[23:16] bit width of the performance monitoring counter
IA32_PMCi[COUNTERWIDTH-1:32] ← EDX[COUNTERWIDTH-33:0];
IA32_PMCi[31:0] ← EAX[31:0];
EDX[63:COUNTERWIDTH] are reserved
```

## 18.3 PERFORMANCE MONITORING (INTEL® CORE™ PROCESSORS AND INTEL® XEON® PROCESSORS)

### 18.3.1 Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Nehalem

Intel Core i7 processor family<sup>2</sup> supports architectural performance monitoring capability with version ID 3 (see Section 18.2.3) and a host of non-architectural monitoring capabilities. The Intel Core i7 processor family is based on Intel® microarchitecture code name Nehalem, and provides four general-purpose performance counters (IA32\_PMC0, IA32\_PMC1, IA32\_PMC2, IA32\_PMC3) and three fixed-function performance counters (IA32\_FIXED\_CTR0, IA32\_FIXED\_CTR1, IA32\_FIXED\_CTR2) in the processor core.

Non-architectural performance monitoring in Intel Core i7 processor family uses the IA32\_PERFVTSELx MSR to configure a set of non-architecture performance monitoring events to be counted by the corresponding general-purpose performance counter. The list of non-architectural performance monitoring events is listed in Table 19-29. Non-architectural performance monitoring events fall into two broad categories:

- Performance monitoring events in the processor core: These include many events that are similar to performance monitoring events available to processor based on Intel Core microarchitecture. Additionally, there are several enhancements in the performance monitoring capability for detecting microarchitectural conditions in the processor core or in the interaction of the processor core to the off-core sub-systems in the physical processor package. The off-core sub-systems in the physical processor package is loosely referred to as “uncore”.
- Performance monitoring events in the uncore: The uncore sub-system is shared by more than one processor cores in the physical processor package. It provides additional performance monitoring facility outside of IA32\_PMCx and performance monitoring events that are specific to the uncore sub-system.

Architectural and non-architectural performance monitoring events in Intel Core i7 processor family support thread qualification using bit 21 of IA32\_PERFVTSELx MSR.

The bit fields within each IA32\_PERFVTSELx MSR are defined in Figure 18-6 and described in Section 18.2.1.1 and Section 18.2.3.

2. Intel Xeon processor 5500 series and 3400 series are also based on Intel microarchitecture code name Nehalem; the performance monitoring facilities described in this section generally also apply.

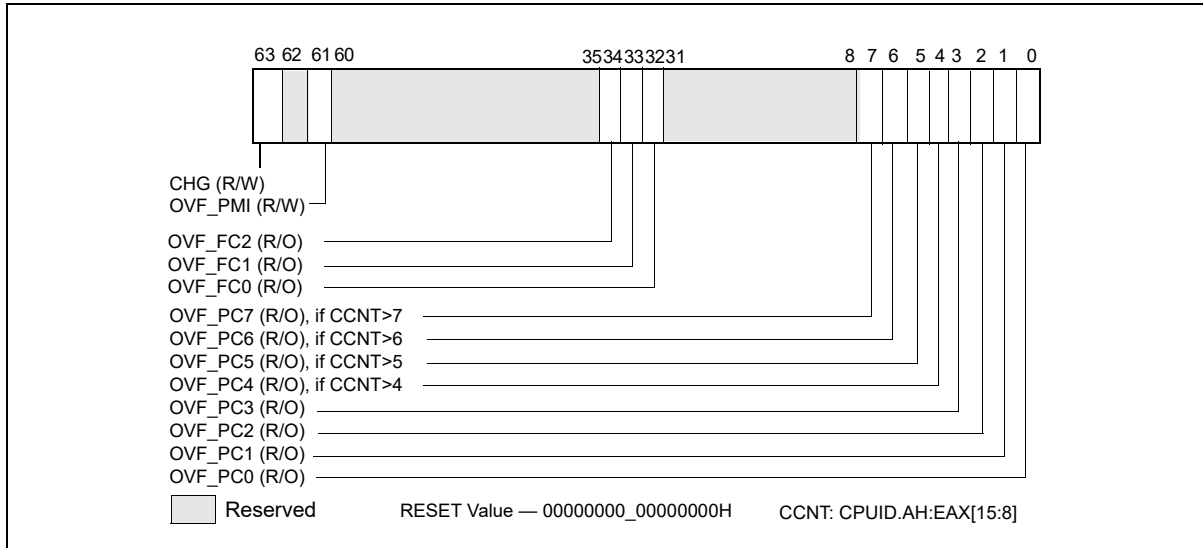


Figure 18-14. IA32\_PERF\_GLOBAL\_STATUS MSR

### 18.3.1.1 Enhancements of Performance Monitoring in the Processor Core

The notable enhancements in the monitoring of performance events in the processor core include:


- Four general purpose performance counters, IA32\_PMCx, associated counter configuration MSRs, IA32\_PERFEVTSELx, and global counter control MSR supporting simplified control of four counters. Each of the four performance counter can support processor event based sampling (PEBS) and thread-qualification of architectural and non-architectural performance events. Width of IA32\_PMCx supported by hardware has been increased. The width of counter reported by CPUID.0AH:EAX[23:16] is 48 bits. The PEBS facility in Intel micro-architecture code name Nehalem has been enhanced to include new data format to capture additional information, such as load latency.
- Load latency sampling facility. Average latency of memory load operation can be sampled using load-latency facility in processors based on Intel microarchitecture code name Nehalem. This field measures the load latency from load's first dispatch of till final data writeback from the memory subsystem. The latency is reported for retired demand load operations and in core cycles (it accounts for re-dispatches). This facility is used in conjunction with the PEBS facility.
- Off-core response counting facility. This facility in the processor core allows software to count certain transaction responses between the processor core to sub-systems outside the processor core (uncore). Counting off-core response requires additional event qualification configuration facility in conjunction with IA32\_PERFEVTSELx. Two off-core response MSRs are provided to use in conjunction with specific event codes that must be specified with IA32\_PERFEVTSELx.

#### 18.3.1.1.1 Processor Event Based Sampling (PEBS)

All four general-purpose performance counters, IA32\_PMCx, can be used for PEBS if the performance event supports PEBS. Software uses IA32\_MISC\_ENABLE[7] and IA32\_MISC\_ENABLE[12] to detect whether the performance monitoring facility and PEBS functionality are supported in the processor. The MSR IA32\_PEBS\_ENABLE provides 4 bits that software must use to enable which IA32\_PMCx overflow condition will cause the PEBS record to be captured.

Additionally, the PEBS record is expanded to allow latency information to be captured. The MSR IA32\_PEBS\_ENABLE provides 4 additional bits that software must use to enable latency data recording in the PEBS record upon the respective IA32\_PMCx overflow condition. The layout of IA32\_PEBS\_ENABLE for processors based on Intel microarchitecture code name Nehalem is shown in Figure 18-15.

When a counter is enabled to capture machine state (PEBS\_EN\_PMCx = 1), the processor will write machine state information to a memory buffer specified by software as detailed below. When the counter IA32\_PMCx overflows from maximum count to zero, the PEBS hardware is armed.



**Figure 18-15. Layout of IA32\_PEBS\_ENABLE MSR**

Upon occurrence of the next PEBS event, the PEBS hardware triggers an assist and causes a PEBS record to be written. The format of the PEBS record is indicated by the bit field IA32\_PERF\_CAPABILITIES[11:8] (see Figure 18-63).

The behavior of PEBS assists is reported by IA32\_PERF\_CAPABILITIES[6] (see Figure 18-63). The return instruction pointer (RIP) reported in the PEBS record will point to the instruction after (+1) the instruction that causes the PEBS assist. The machine state reported in the PEBS record is the machine state after the instruction that causes the PEBS assist is retired. For instance, if the instructions:

```
mov eax, [eax] ; causes PEBS assist
```

```
nop
```

are executed, the PEBS record will report the address of the nop, and the value of EAX in the PEBS record will show the value read from memory, not the target address of the read operation.

The PEBS record format is shown in Table 18-3, and each field in the PEBS record is 64 bits long. The PEBS record format, along with debug/store area storage format, does not change regardless of IA-32e mode is active or not. CPUID.01H:ECX.DTES64[bit 2] reports whether the processor's DS storage format support is mode-independent. When set, it uses 64-bit DS storage format.

**Table 18-3. PEBS Record Format for Intel Core i7 Processor Family**

| Byte Offset | Field    | Byte Offset | Field                   |
|-------------|----------|-------------|-------------------------|
| 00H         | R/EFLAGS | 58H         | R9                      |
| 08H         | R/EIP    | 60H         | R10                     |
| 10H         | R/EAX    | 68H         | R11                     |
| 18H         | R/EBX    | 70H         | R12                     |
| 20H         | R/ECX    | 78H         | R13                     |
| 28H         | R/EDX    | 80H         | R14                     |
| 30H         | R/ESI    | 88H         | R15                     |
| 38H         | R/EDI    | 90H         | IA32_PERF_GLOBAL_STATUS |
| 40H         | R/EBP    | 98H         | Data Linear Address     |
| 48H         | R/ESP    | A0H         | Data Source Encoding    |

**Table 18-3. PEBS Record Format for Intel Core i7 Processor Family**

| Byte Offset | Field | Byte Offset | Field                       |
|-------------|-------|-------------|-----------------------------|
| 50H         | R8    | A8H         | Latency value (core cycles) |

In IA-32e mode, the full 64-bit value is written to the register. If the processor is not operating in IA-32e mode, 32-bit value is written to registers with bits 63:32 zeroed. Registers not defined when the processor is not in IA-32e mode are written to zero.

Bytes AFH:90H are enhancement to the PEBS record format. Support for this enhanced PEBS record format is indicated by IA32\_PERF\_CAPABILITIES[11:8] encoding of 0001B.

The value written to bytes 97H:90H is the state of the IA32\_PERF\_GLOBAL\_STATUS register before the PEBS assist occurred. This value is written so software can determine which counters overflowed when this PEBS record was written. Note that this field indicates the overflow status for all counters, regardless of whether they were programmed for PEBS or not.

**Programming PEBS Facility**

Only a subset of non-architectural performance events in the processor support PEBS. The subset of precise events are listed in Table 18-68. In addition to using IA32\_PERFEVTSELx to specify event unit/mask settings and setting the EN\_PMCx bit in the IA32\_PEBS\_ENABLE register for the respective counter, the software must also initialize the DS\_BUFFER\_MANAGEMENT\_AREA data structure in memory to support capturing PEBS records for precise events.

**NOTE**

PEBS events are only valid when the following fields of IA32\_PERFEVTSELx are all zero: AnyThread, Edge, Invert, CMask.

The beginning linear address of the DS\_BUFFER\_MANAGEMENT\_AREA data structure must be programmed into the IA32\_DS\_AREA register. The layout of the DS\_BUFFER\_MANAGEMENT\_AREA is shown in Figure 18-16.

- **PEBS Buffer Base:** This field is programmed with the linear address of the first byte of the PEBS buffer allocated by software. The processor reads this field to determine the base address of the PEBS buffer. Software should allocate this memory from the non-paged pool.
- **PEBS Index:** This field is initially programmed with the same value as the PEBS Buffer Base field, or the beginning linear address of the PEBS buffer. The processor reads this field to determine the location of the next PEBS record to write to. After a PEBS record has been written, the processor also updates this field with the address of the next PEBS record to be written. The figure above illustrates the state of PEBS Index after the first PEBS record is written.
- **PEBS Absolute Maximum:** This field represents the absolute address of the maximum length of the allocated PEBS buffer plus the starting address of the PEBS buffer. The processor will not write any PEBS record beyond the end of PEBS buffer, when PEBS Index equals PEBS Absolute Maximum. No signaling is generated when PEBS buffer is full. Software must reset the PEBS Index field to the beginning of the PEBS buffer address to continue capturing PEBS records.




Figure 18-16. PEBS Programming Environment

- PEBS Interrupt Threshold:** This field specifies the threshold value to trigger a performance interrupt and notify software that the PEBS buffer is nearly full. This field is programmed with the linear address of the first byte of the PEBS record within the PEBS buffer that represents the threshold record. After the processor writes a PEBS record and updates **PEBS Index**, if the **PEBS Index** reaches the threshold value of this field, the processor will generate a performance interrupt. This is the same interrupt that is generated by a performance counter overflow, as programmed in the Performance Monitoring Counters vector in the Local Vector Table of the Local APIC. When a performance interrupt due to PEBS buffer full is generated, the `IA32_PERF_GLOBAL_STATUS.PEBS_Ovf` bit will be set.
- PEBS CounterX Reset:** This field allows software to set up PEBS counter overflow condition to occur at a rate useful for profiling workload, thereby generating multiple PEBS records to facilitate characterizing the profile the execution of test code. After each PEBS record is written, the processor checks each counter to see if it overflowed and was enabled for PEBS (the corresponding bit in `IA32_PEBS_ENABLED` was set). If these conditions are met, then the reset value for each overflowed counter is loaded from the DS Buffer Management Area. For example, if counter `IA32_PMC0` caused a PEBS record to be written, then the value of "PEBS Counter 0 Reset" would be written to counter `IA32_PMC0`. If a counter is not enabled for PEBS, its value will not be modified by the PEBS assist.

#### Performance Counter Prioritization

Performance monitoring interrupts are triggered by a counter transitioning from maximum count to zero (assuming `IA32_PerfEvtSelX.INT` is set). This same transition will cause PEBS hardware to arm, but not trigger. PEBS hardware triggers upon detection of the first PEBS event after the PEBS hardware has been armed (a 0 to 1 transition of the counter). At this point, a PEBS assist will be undertaken by the processor.

Performance counters (fixed and general-purpose) are prioritized in index order. That is, counter IA32\_PMC0 takes precedence over all other counters. Counter IA32\_PMC1 takes precedence over counters IA32\_PMC2 and IA32\_PMC3, and so on. This means that if simultaneous overflows or PEBS assists occur, the appropriate action will be taken for the highest priority performance counter. For example, if IA32\_PMC1 cause an overflow interrupt and IA32\_PMC2 causes an PEBS assist simultaneously, then the overflow interrupt will be serviced first.

The PEBS threshold interrupt is triggered by the PEBS assist, and is by definition prioritized lower than the PEBS assist. Hardware will not generate separate interrupts for each counter that simultaneously overflows. General-purpose performance counters are prioritized over fixed counters.

If a counter is programmed with a precise (PEBS-enabled) event and programmed to generate a counter overflow interrupt, the PEBS assist is serviced before the counter overflow interrupt is serviced. If in addition the PEBS interrupt threshold is met, the

threshold interrupt is generated after the PEBS assist completes, followed by the counter overflow interrupt (two separate interrupts are generated).

Uncore counters may be programmed to interrupt one or more processor cores (see Section 18.3.1.2). It is possible for interrupts posted from the uncore facility to occur coincident with counter overflow interrupts from the processor core. Software must check core and uncore status registers to determine the exact origin of counter overflow interrupts.

### 18.3.1.1.2 Load Latency Performance Monitoring Facility

The load latency facility provides software a means to characterize the average load latency to different levels of cache/memory hierarchy. This facility requires processor supporting enhanced PEBS record format in the PEBS buffer, see Table 18-3. This field measures the load latency from load's first dispatch of till final data writeback from the memory subsystem. The latency is reported for retired demand load operations and in core cycles (it accounts for re-dispatches).

To use this feature software must assure:

- One of the IA32\_PERFEVTSELx MSR is programmed to specify the event unit MEM\_INST\_RETIRED, and the LATENCY\_ABOVE\_THRESHOLD event mask must be specified (IA32\_PerfEvtSelX[15:0] = 100H). The corresponding counter IA32\_PMCx will accumulate event counts for architecturally visible loads which exceed the programmed latency threshold specified separately in a MSR. Stores are ignored when this event is programmed. The CMASK or INV fields of the IA32\_PerfEvtSelX register used for counting load latency must be 0. Writing other values will result in undefined behavior.
- The MSR\_PEBS\_LD\_LAT\_THRESHOLD MSR is programmed with the desired latency threshold in core clock cycles. Loads with latencies greater than this value are eligible for counting and latency data reporting. The minimum value that may be programmed in this register is 3 (the minimum detectable load latency is 4 core clock cycles).
- The PEBS enable bit in the IA32\_PEBS\_ENABLE register is set for the corresponding IA32\_PMCx counter register. This means that both the PEBS\_EN\_CTRX and LL\_EN\_CTRX bits must be set for the counter(s) of interest. For example, to enable load latency on counter IA32\_PMC0, the IA32\_PEBS\_ENABLE register must be programmed with the 64-bit value 00000001\_00000001H.

When the load-latency facility is enabled, load operations are randomly selected by hardware and tagged to carry information related to data source locality and latency. Latency and data source information of tagged loads are updated internally.

When a PEBS assist occurs, the last update of latency and data source information are captured by the assist and written as part of the PEBS record. The PEBS sample after value (SAV), specified in PEBS CounterX Reset, operates orthogonally to the tagging mechanism. Loads are randomly tagged to collect latency data. The SAV controls the number of tagged loads with latency information that will be written into the PEBS record field by the PEBS assists. The load latency data written to the PEBS record will be for the last tagged load operation which retired just before the PEBS assist was invoked.

The load-latency information written into a PEBS record (see Table 18-3, bytes AFH:98H) consists of:

- **Data Linear Address:** This is the linear address of the target of the load operation.
- **Latency Value:** This is the elapsed cycles of the tagged load operation between dispatch to GO, measured in processor core clock domain.

- **Data Source:** The encoded value indicates the origin of the data obtained by the load instruction. The encoding is shown in Table 18-4. In the descriptions local memory refers to system memory physically attached to a processor package, and remote memory referrals to system memory physically attached to another processor package.


**Table 18-4. Data Source Encoding for Load Latency Record**

| Encoding         | Description   |
|------------------|---|
| 00H              | Unknown L3 cache miss   |
| 01H              | Minimal latency core cache hit. This request was satisfied by the L1 data cache.  |
| 02H              | Pending core cache HIT. Outstanding core cache miss to same cache-line address was already underway.  |
| 03H              | This data request was satisfied by the L2.  |
| 04H              | L3 HIT. Local or Remote home requests that hit L3 cache in the uncore with no coherency actions required (snooping).  |
| 05H              | L3 HIT. Local or Remote home requests that hit the L3 cache and was serviced by another processor core with a cross core snoop where no modified copies were found. (clean).                      |
| 06H              | L3 HIT. Local or Remote home requests that hit the L3 cache and was serviced by another processor core with a cross core snoop where modified copies were found. (HITM).                          |
| 07H <sup>1</sup> | Reserved/LLC Snoop HitM. Local or Remote home requests that hit the last level cache and was serviced by another core with a cross core snoop where modified copies found                         |
| 08H              | L3 MISS. Local homed requests that missed the L3 cache and was serviced by forwarded data following a cross package snoop where no modified copies found. (Remote home requests are not counted). |
| 09H              | Reserved  |
| 0AH              | L3 MISS. Local home requests that missed the L3 cache and was serviced by local DRAM (go to shared state).  |
| 0BH              | L3 MISS. Remote home requests that missed the L3 cache and was serviced by remote DRAM (go to shared state).  |
| 0CH              | L3 MISS. Local home requests that missed the L3 cache and was serviced by local DRAM (go to exclusive state).   |
| 0DH              | L3 MISS. Remote home requests that missed the L3 cache and was serviced by remote DRAM (go to exclusive state).   |
| 0EH              | I/O, Request of input/output operation  |
| 0FH              | The request was to un-cacheable memory.   |

**NOTES:**

1. Bit 7 is supported only for processor with CPUID DisplayFamily\_DisplayModel signature of 06\_2A, and 06\_2E; otherwise it is reserved.

The layout of MSR\_PEBS\_LD\_LAT\_THRESHOLD is shown in Figure 18-17.



**Figure 18-17. Layout of MSR\_PEBS\_LD\_LAT MSR**

Bits 15:0 specifies the threshold load latency in core clock cycles. Performance events with latencies greater than this value are counted in IA32\_PMCx and their latency information is reported in the PEBS record. Otherwise, they are ignored. The minimum value that may be programmed in this field is 3.

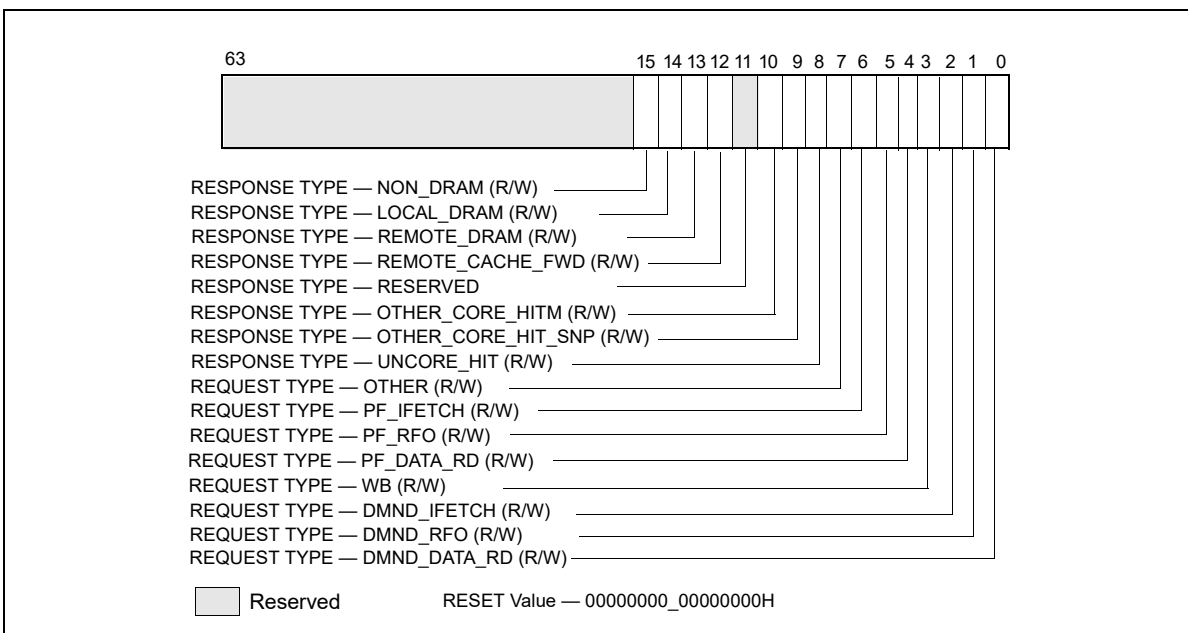
### 18.3.1.1.3 Off-core Response Performance Monitoring in the Processor Core

Programming a performance event using the off-core response facility can choose any of the four IA32\_PERFEVTSELx MSR with specific event codes and predefine mask bit value. Each event code for off-core response monitoring requires programming an associated configuration MSR, MSR\_OFFCORE\_RSP\_0. There is only one off-core response configuration MSR. Table 18-5 lists the event code, mask value and additional off-core configuration MSR that must be programmed to count off-core response events using IA32\_PMCx.

**Table 18-5. Off-Core Response Event Encoding**

| Event code in IA32_PERFEVTSELx | Mask Value in IA32_PERFEVTSELx | Required Off-core Response MSR   |
|--------------------------------|--------------------------------|----------------------------------|
| B7H                            | 01H                            | MSR_OFFCORE_RSP_0 (address 1A6H) |

The layout of MSR\_OFFCORE\_RSP\_0 is shown in Figure 18-18. Bits 7:0 specifies the request type of a transaction request to the uncore. Bits 15:8 specifies the response of the uncore subsystem.



**Figure 18-18. Layout of MSR\_OFFCORE\_RSP\_0 and MSR\_OFFCORE\_RSP\_1 to Configure Off-core Response Events**

**Table 18-6. MSR\_OFFCORE\_RSP\_0 and MSR\_OFFCORE\_RSP\_1 Bit Field Definition**

| Bit Name     | Offset | Description   |
|--------------|--------|---|
| DMND_DATA_RD | 0      | (R/W). Counts the number of demand and DCU prefetch data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count L2 data read prefetches or instruction fetches. |
| DMND_RFO     | 1      | (R/W). Counts the number of demand and DCU prefetch reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO.   |
| DMND_IFETCH  | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.  |
| WB           | 3      | (R/W). Counts the number of writeback (modified to exclusive) transactions.   |
| PF_DATA_RD   | 4      | (R/W). Counts the number of data cacheline reads generated by L2 prefetchers.   |
| PF_RFO       | 5      | (R/W). Counts the number of RFO requests generated by L2 prefetchers.   |



**Table 18-6. MSR\_OFFCORE\_RSP\_0 and MSR\_OFFCORE\_RSP\_1 Bit Field Definition (Contd.)**

| Bit Name           | Offset | Description   |
|--------------------|--------|---|
| PF_IFETCH          | 6      | (R/W). Counts the number of code reads generated by L2 prefetchers.   |
| OTHER              | 7      | (R/W). Counts one of the following transaction types, including L3 invalidate, I/O, full or partial writes, WC or non-temporal stores, CLFLUSH, Fences, lock, unlock, split lock.                       |
| UNCORE_HIT         | 8      | (R/W). L3 Hit: local or remote home requests that hit L3 cache in the uncore with no coherency actions required (snooping).   |
| OTHER_CORE_HIT_SNP | 9      | (R/W). L3 Hit: local or remote home requests that hit L3 cache in the uncore and was serviced by another core with a cross core snoop where no modified copies were found (clean).                      |
| OTHER_CORE_HIT_TM  | 10     | (R/W). L3 Hit: local or remote home requests that hit L3 cache in the uncore and was serviced by another core with a cross core snoop where modified copies were found (HITM).                          |
| Reserved           | 11     | Reserved  |
| REMOTE_CACHE_FWD   | 12     | (R/W). L3 Miss: local homed requests that missed the L3 cache and was serviced by forwarded data following a cross package snoop where no modified copies found. (Remote home requests are not counted) |
| REMOTE_DRAM        | 13     | (R/W). L3 Miss: remote home requests that missed the L3 cache and were serviced by remote DRAM.   |
| LOCAL_DRAM         | 14     | (R/W). L3 Miss: local home requests that missed the L3 cache and were serviced by local DRAM.   |
| NON_DRAM           | 15     | (R/W). Non-DRAM requests that were serviced by IOH.   |

### 18.3.1.2 Performance Monitoring Facility in the Uncore

The “uncore” in Intel microarchitecture code name Nehalem refers to subsystems in the physical processor package that are shared by multiple processor cores. Some of the sub-systems in the uncore include the L3 cache, Intel QuickPath Interconnect link logic, and integrated memory controller. The performance monitoring facilities inside the uncore operates in the same clock domain as the uncore (U-clock domain), which is usually different from the processor core clock domain. The uncore performance monitoring facilities described in this section apply to Intel Xeon processor 5500 series and processors with the following CPUID signatures: 06\_1AH, 06\_1EH, 06\_1FH (see Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*). An overview of the uncore performance monitoring facilities is described separately.

The performance monitoring facilities available in the U-clock domain consist of:

- Eight General-purpose counters (MSR\_UNCORE\_PerfCntr0 through MSR\_UNCORE\_PerfCntr7). The counters are 48 bits wide. Each counter is associated with a configuration MSR, MSR\_UNCORE\_PerfEvtSelx, to specify event code, event mask and other event qualification fields. A set of global uncore performance counter enabling/overflow/status control MSRs are also provided for software.
- Performance monitoring in the uncore provides an address/opcode match MSR that provides event qualification control based on address value or QPI command opcode.
- One fixed-function counter, MSR\_UNCORE\_FixedCntr0. The fixed-function uncore counter increments at the rate of the U-clock when enabled.

The frequency of the uncore clock domain can be determined from the uncore clock ratio which is available in the PCI configuration space register at offset COH under device number 0 and Function 0.

#### 18.3.1.2.1 Uncore Performance Monitoring Management Facility

MSR\_UNCORE\_PERF\_GLOBAL\_CTRL provides bit fields to enable/disable general-purpose and fixed-function counters in the uncore. Figure 18-19 shows the layout of MSR\_UNCORE\_PERF\_GLOBAL\_CTRL for an uncore that is shared by four processor cores in a physical package.

- EN\_PCn (bit n, n = 0, 7): When set, enables counting for the general-purpose uncore counter MSR\_UNCORE\_PerfCntr n.
- EN\_FC0 (bit 32): When set, enables counting for the fixed-function uncore counter MSR\_UNCORE\_FixedCntr0.

- EN\_PMI\_COREn (bit n, n = 0, 3 if four cores are present): When set, processor core n is programmed to receive an interrupt signal from any interrupt enabled uncore counter. PMI delivery due to an uncore counter overflow is enabled by setting IA32\_DEBUGCTL.Offcore\_PMI\_EN to 1.
- PMI\_FRZ (bit 63): When set, all U-clock uncore counters are disabled when any one of them signals a performance interrupt. Software must explicitly re-enable the counter by setting the enable bits in MSR\_UNCORE\_PERF\_GLOBAL\_CTRL upon exit from the ISR.




Figure 18-19. Layout of MSR\_UNCORE\_PERF\_GLOBAL\_CTRL MSR

MSR\_UNCORE\_PERF\_GLOBAL\_STATUS provides overflow status of the U-clock performance counters in the uncore. This is a read-only register. If an overflow status bit is set the corresponding counter has overflowed. The register provides a condition change bit (bit 63) which can be quickly checked by software to determine if a significant change has occurred since the last time the condition change status was cleared. Figure 18-20 shows the layout of MSR\_UNCORE\_PERF\_GLOBAL\_STATUS.

- OVF\_PCn (bit n, n = 0, 7): When set, indicates general-purpose uncore counter MSR\_UNCORE\_PerfCntr n has overflowed.
- OVF\_FC0 (bit 32): When set, indicates the fixed-function uncore counter MSR\_UNCORE\_FixedCntr0 has overflowed.
- OVF\_PMI (bit 61): When set indicates that an uncore counter overflowed and generated an interrupt request.
- CHG (bit 63): When set indicates that at least one status bit in MSR\_UNCORE\_PERF\_GLOBAL\_STATUS register has changed state.

MSR\_UNCORE\_PERF\_GLOBAL\_OVF\_CTRL allows software to clear the status bits in the UNCORE\_PERF\_GLOBAL\_STATUS register. This is a write-only register, and individual status bits in the global status register are cleared by writing a binary one to the corresponding bit in this register. Writing zero to any bit position in this register has no effect on the uncore PMU hardware.




Figure 18-20. Layout of MSR\_UNCORE\_PERF\_GLOBAL\_STATUS MSR

Figure 18-21 shows the layout of MSR\_UNCORE\_PERF\_GLOBAL\_OVF\_CTRL.




Figure 18-21. Layout of MSR\_UNCORE\_PERF\_GLOBAL\_OVF\_CTRL MSR

- CLR\_OVF\_PCn (bit n, n = 0, 7): Set this bit to clear the overflow status for general-purpose uncore counter MSR\_UNCORE\_PerfCntr n. Writing a value other than 1 is ignored.
- CLR\_OVF\_FC0 (bit 32): Set this bit to clear the overflow status for the fixed-function uncore counter MSR\_UNCORE\_FixedCntr0. Writing a value other than 1 is ignored.
- CLR\_OVF\_PMI (bit 61): Set this bit to clear the OVF\_PMI flag in MSR\_UNCORE\_PERF\_GLOBAL\_STATUS. Writing a value other than 1 is ignored.
- CLR\_CHG (bit 63): Set this bit to clear the CHG flag in MSR\_UNCORE\_PERF\_GLOBAL\_STATUS register. Writing a value other than 1 is ignored.

### 18.3.1.2.2 Uncore Performance Event Configuration Facility

MSR\_UNCORE\_PerfEvtSel0 through MSR\_UNCORE\_PerfEvtSel7 are used to select performance event and configure the counting behavior of the respective uncore performance counter. Each uncore PerfEvtSel MSR is paired with an uncore performance counter. Each uncore counter must be locally configured using the corresponding MSR\_UNCORE\_PerfEvtSelx and counting must be enabled using the respective EN\_PCx bit in MSR\_UNCORE\_PERF\_GLOBAL\_CTRL. Figure 18-22 shows the layout of MSR\_UNCORE\_PERFEVTSELx.

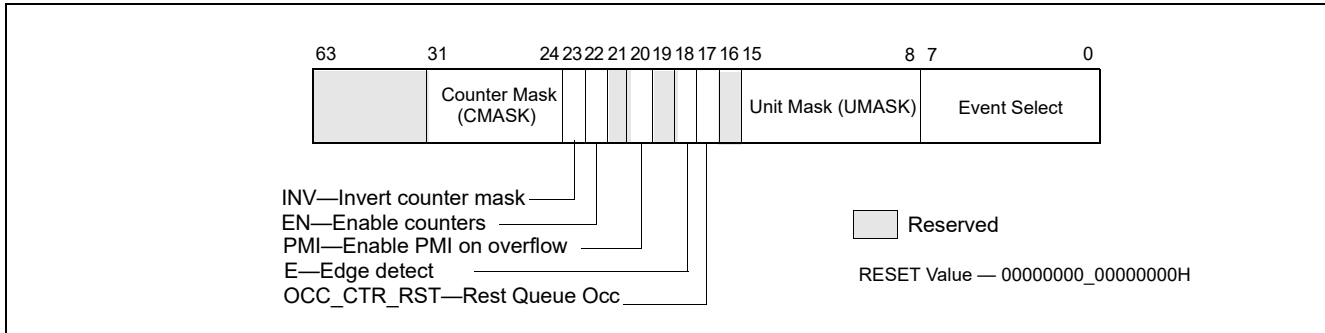


Figure 18-22. Layout of MSR\_UNCORE\_PERFEVTSELx MSRs

- Event Select (bits 7:0): Selects the event logic unit used to detect uncore events.
- Unit Mask (bits 15:8) : Condition qualifiers for the event selection logic specified in the Event Select field.
- OCC\_CTR\_RST (bit17): When set causes the queue occupancy counter associated with this event to be cleared (zeroed). Writing a zero to this bit will be ignored. It will always read as a zero.
- Edge Detect (bit 18): When set causes the counter to increment when a deasserted to asserted transition occurs for the conditions that can be expressed by any of the fields in this register.
- PMI (bit 20): When set, the uncore will generate an interrupt request when this counter overflowed. This request will be routed to the logical processors as enabled in the PMI enable bits (EN\_PMI\_COREx) in the register MSR\_UNCORE\_PERF\_GLOBAL\_CTRL.
- EN (bit 22): When clear, this counter is locally disabled. When set, this counter is locally enabled and counting starts when the corresponding EN\_PCx bit in MSR\_UNCORE\_PERF\_GLOBAL\_CTRL is set.
- INV (bit 23): When clear, the Counter Mask field is interpreted as greater than or equal to. When set, the Counter Mask field is interpreted as less than.
- Counter Mask (bits 31:24): When this field is clear, it has no effect on counting. When set to a value other than zero, the logical processor compares this field to the event counts on each core clock cycle. If INV is clear and the event counts are greater than or equal to this field, the counter is incremented by one. If INV is set and the event counts are less than this field, the counter is incremented by one. Otherwise the counter is not incremented.

Figure 18-23 shows the layout of MSR\_UNCORE\_FIXED\_CTR\_CTRL.




Figure 18-23. Layout of MSR\_UNCORE\_FIXED\_CTR\_CTRL MSR

- EN (bit 0): When clear, the uncore fixed-function counter is locally disabled. When set, it is locally enabled and counting starts when the EN\_FC0 bit in MSR\_UNCORE\_PERF\_GLOBAL\_CTRL is set.
- PMI (bit 2): When set, the uncore will generate an interrupt request when the uncore fixed-function counter overflowed. This request will be routed to the logical processors as enabled in the PMI enable bits (EN\_PMI\_COREx) in the register MSR\_UNCORE\_PERF\_GLOBAL\_CTRL.

Both the general-purpose counters (MSR\_UNCORE\_PerfCnt) and the fixed-function counter (MSR\_UNCORE\_FixedCnt0) are 48 bits wide. They support both counting and interrupt based sampling usages. The event logic unit can filter event counts to specific regions of code or transaction types incoming to the home node logic.

### 18.3.1.2.3 Uncore Address/Opcode Match MSR

The Event Select field [7:0] of MSR\_UNCORE\_PERFEVTSELx is used to select different uncore event logic unit. When the event "ADDR\_OPCODE\_MATCH" is selected in the Event Select field, software can filter uncore performance events according to transaction address and certain transaction responses. The address filter and transaction response filtering requires the use of MSR\_UNCORE\_ADDR\_OPCODE\_MATCH register. The layout is shown in Figure 18-24.

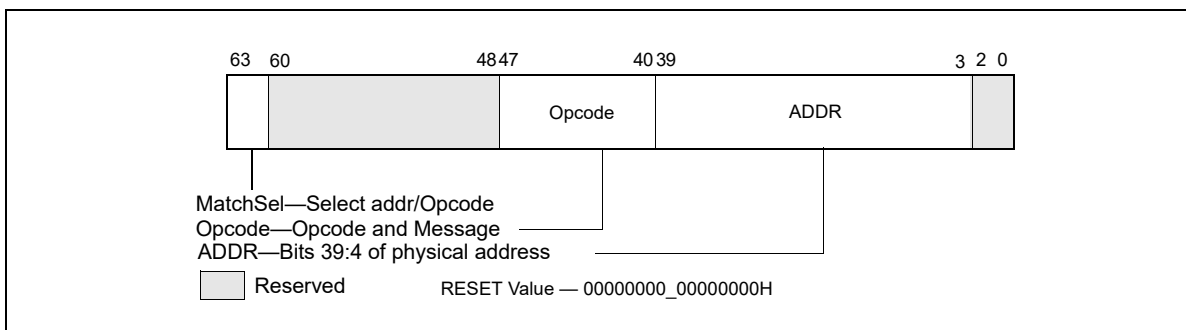


Figure 18-24. Layout of MSR\_UNCORE\_ADDR\_OPCODE\_MATCH MSR

- Addr (bits 39:3): The physical address to match if "MatchSel" field is set to select address match. The uncore performance counter will increment if the lowest 40-bit incoming physical address (excluding bits 2:0) for a transaction request matches bits 39:3.
- Opcode (bits 47:40) : Bits 47:40 allow software to filter uncore transactions based on QPI link message class/packed header opcode. These bits are consists two sub-fields:
  - Bits 43:40 specify the QPI packet header opcode.
  - Bits 47:44 specify the QPI message classes.

Table 18-7 lists the encodings supported in the opcode field.

**Table 18-7. Opcode Field Encoding for MSR\_UNCORE\_ADDR\_OPCODE\_MATCH**

| Opcode [43:40] | QPI Message Class               |                                   |                                  |
|----------------|---------------------------------|-----------------------------------|----------------------------------|
|                | Home Request<br>[47:44] = 0000B | Snoop Response<br>[47:44] = 0001B | Data Response<br>[47:44] = 1110B |
|                |                                 | 1                                 |                                  |
| DMND_IFETCH    | 2                               | 2                                 |                                  |
| WB             | 3                               | 3                                 |                                  |
| PF_DATA_RD     | 4                               | 4                                 |                                  |
| PF_RFO         | 5                               | 5                                 |                                  |
| PF_IFETCH      | 6                               | 6                                 |                                  |
| OTHER          | 7                               | 7                                 |                                  |
| NON_DRAM       | 15                              | 15                                |                                  |

- MatchSel (bits 63:61): Software specifies the match criteria according to the following encoding:
  - 000B: Disable addr\_opcode match hardware.
  - 100B: Count if only the address field matches.
  - 010B: Count if only the opcode field matches.
  - 110B: Count if either opcode field matches or the address field matches.
  - 001B: Count only if both opcode and address field match.
  - Other encoding are reserved.

### 18.3.1.3 Intel® Xeon® Processor 7500 Series Performance Monitoring Facility

The performance monitoring facility in the processor core of Intel® Xeon® processor 7500 series are the same as those supported in Intel Xeon processor 5500 series. The uncore subsystem in Intel Xeon processor 7500 series are significantly different. The uncore performance monitoring facility consist of many distributed units associated with individual logic control units (referred to as boxes) within the uncore subsystem. A high level block diagram of the various box units of the uncore is shown in Figure 18-25.

Uncore PMUs are programmed via MSR interfaces. Each of the distributed uncore PMU units have several general-purpose counters. Each counter requires an associated event select MSR, and may require additional MSRs to configure sub-event conditions. The uncore PMU MSRs associated with each box can be categorized based on its functional scope: per-counter, per-box, or global across the uncore. The number counters available in each box type are different. Each box generally provides a set of MSRs to enable/disable, check status/overflow of multiple counters within each box.

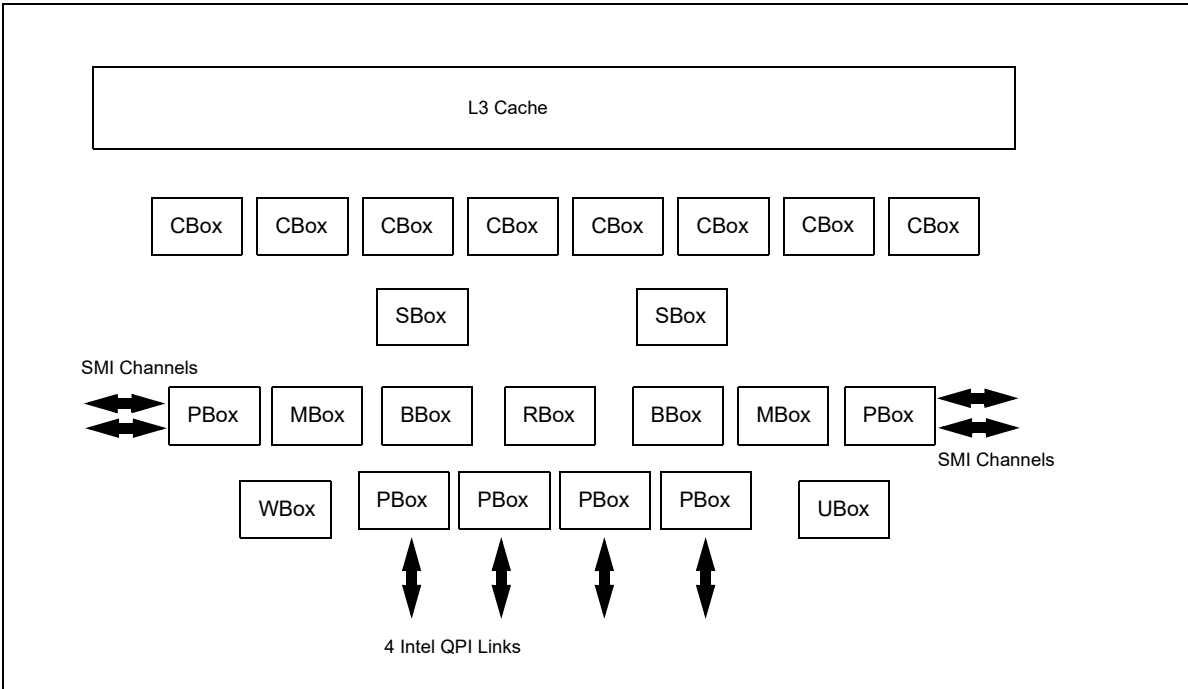


Figure 18-25. Distributed Units of the Uncore of Intel® Xeon® Processor 7500 Series

Table 18-8 summarizes the number MSRs for uncore PMU for each box.

Table 18-8. Uncore PMU MSR Summary

| Box   | # of Boxes | Counters per Box         | Counter Width | General Purpose | Global Enable | Sub-control MSRs |
|-------|------------|--------------------------|---------------|-----------------|---------------|------------------|
| C-Box | 8          | 6                        | 48            | Yes             | per-box       | None             |
| S-Box | 2          | 4                        | 48            | Yes             | per-box       | Match/Mask       |
| B-Box | 2          | 4                        | 48            | Yes             | per-box       | Match/Mask       |
| M-Box | 2          | 6                        | 48            | Yes             | per-box       | Yes              |
| R-Box | 1          | 16 ( 2 port, 8 per port) | 48            | Yes             | per-box       | Yes              |
| W-Box | 1          | 4                        | 48            | Yes             | per-box       | None             |
|       |            | 1                        | 48            | No              | per-box       | None             |
| U-Box | 1          | 1                        | 48            | Yes             | uncore        | None             |

The W-Box provides 4 general-purpose counters, each requiring an event select configuration MSR, similar to the general-purpose counters in other boxes. There is also a fixed-function counter that increments clockticks in the uncore clock domain.

For C,S,B,M,R, and W boxes, each box provides an MSR to enable/disable counting, configuring PMI of multiple counters within the same box, this is somewhat similar the “global control” programming interface, IA32\_PERF\_GLOBAL\_CTRL, offered in the core PMU. Similarly status information and counter overflow control for multiple counters within the same box are also provided in C,S,B,M,R, and W boxes.

In the U-Box, MSR\_U\_PMON\_GLOBAL\_CTL provides overall uncore PMU enable/disable and PMI configuration control. The scope of status information in the U-box is at per-box granularity, in contrast to the per-box status information MSR (in the C,S,B,M,R, and W boxes) providing status information of individual counter overflow. The difference in scope also apply to the overflow control MSR in the U-Box versus those in the other Boxes.

The individual MSRs that provide uncore PMU interfaces are listed in Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*, Table 2-16 under the general naming style of MSR\_%box#%\_PMON\_%scope\_function%, where %box#% designates the type of box and zero-based index if there are more than one box of the same type, %scope\_function% follows the examples below:

- Multi-counter enabling MSRs: MSR\_U\_PMON\_GLOBAL\_CTL, MSR\_S0\_PMON\_BOX\_CTL, MSR\_C7\_PMON\_BOX\_CTL, etc.
- Multi-counter status MSRs: MSR\_U\_PMON\_GLOBAL\_STATUS, MSR\_S0\_PMON\_BOX\_STATUS, MSR\_C7\_PMON\_BOX\_STATUS, etc.
- Multi-counter overflow control MSRs: MSR\_U\_PMON\_GLOBAL\_OVF\_CTL, MSR\_S0\_PMON\_BOX\_OVF\_CTL, MSR\_C7\_PMON\_BOX\_OVF\_CTL, etc.
- Performance counters MSRs: the scope is implicitly per counter, e.g. MSR\_U\_PMON\_CTR, MSR\_S0\_PMON\_CTR0, MSR\_C7\_PMON\_CTR5, etc.
- Event select MSRs: the scope is implicitly per counter, e.g. MSR\_U\_PMON\_EVNT\_SEL, MSR\_S0\_PMON\_EVNT\_SEL0, MSR\_C7\_PMON\_EVNT\_SEL5, etc.
- Sub-control MSRs: the scope is implicitly per-box granularity, e.g. MSR\_M0\_PMON\_TIMESTAMP, MSR\_R0\_PMON\_IPERFO\_P1, MSR\_S1\_PMON\_MATCH.

Details of uncore PMU MSR bit field definitions can be found in a separate document “Intel Xeon Processor 7500 Series Uncore Performance Monitoring Guide”.

### 18.3.2 Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Westmere

All of the performance monitoring programming interfaces (architectural and non-architectural core PMU facilities, and uncore PMU) described in Section 18.6.3 also apply to processors based on Intel® microarchitecture code name Westmere.

Table 18-5 describes a non-architectural performance monitoring event (event code 0B7H) and associated MSR\_OFFCORE\_RSP\_0 (address 1A6H) in the core PMU. This event and a second functionally equivalent offcore response event using event code 0BBH and MSR\_OFFCORE\_RSP\_1 (address 1A7H) are supported in processors based on Intel microarchitecture code name Westmere. The event code and event mask definitions of Non-architectural performance monitoring events are listed in Table 19-29.

The load latency facility is the same as described in Section 18.3.1.1.2, but added enhancement to provide more information in the data source encoding field of each load latency record. The additional information relates to STLB\_MISS and LOCK, see Table 18-13.

### 18.3.3 Intel® Xeon® Processor E7 Family Performance Monitoring Facility

The performance monitoring facility in the processor core of the Intel® Xeon® processor E7 family is the same as those supported in the Intel Xeon processor 5600 series<sup>3</sup>. The uncore subsystem in the Intel Xeon processor E7 family is similar to those of the Intel Xeon processor 7500 series. The high level construction of the uncore subsystem is similar to that shown in Figure 18-25, with the additional capability that up to 10 C-Box units are supported.

---

3. Exceptions are indicated for event code 0FH in Table 19-21; and valid bits of data source encoding field of each load latency record is limited to bits 5:4 of Table 18-13.



Table 18-9 summarizes the number MSRs for uncore PMU for each box.

**Table 18-9. Uncore PMU MSR Summary for Intel® Xeon® Processor E7 Family**

| Box   | # of Boxes | Counters per Box         | Counter Width | General Purpose | Global Enable | Sub-control MSRs |
|-------|------------|--------------------------|---------------|-----------------|---------------|------------------|
| C-Box | 10         | 6                        | 48            | Yes             | per-box       | None             |
| S-Box | 2          | 4                        | 48            | Yes             | per-box       | Match/Mask       |
| B-Box | 2          | 4                        | 48            | Yes             | per-box       | Match/Mask       |
| M-Box | 2          | 6                        | 48            | Yes             | per-box       | Yes              |
| R-Box | 1          | 16 ( 2 port, 8 per port) | 48            | Yes             | per-box       | Yes              |
| W-Box | 1          | 4                        | 48            | Yes             | per-box       | None             |
|       |            | 1                        | 48            | No              | per-box       | None             |
| U-Box | 1          | 1                        | 48            | Yes             | uncore        | None             |

Details of the uncore performance monitoring facility of Intel Xeon Processor E7 family is available in the “Intel® Xeon® Processor E7 Uncore Performance Monitoring Programming Reference Manual”.

### 18.3.4 Performance Monitoring for Processors Based on Intel® Microarchitecture Code Name Sandy Bridge

Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series, and Intel® Xeon® processor E3-1200 family are based on Intel microarchitecture code name Sandy Bridge; this section describes the performance monitoring facilities provided in the processor core. The core PMU supports architectural performance monitoring capability with version ID 3 (see Section 18.2.3) and a host of non-architectural monitoring capabilities.

Architectural performance monitoring version 3 capabilities are described in Section 18.2.3.

The core PMU’s capability is similar to those described in Section 18.3.1.1 and Section 18.6.3, with some differences and enhancements relative to Intel microarchitecture code name Westmere summarized in Table 18-10.

**Table 18-10. Core PMU Comparison**


| Box  | Intel® microarchitecture code name Sandy Bridge  | Intel® microarchitecture code name Westmere  | Comment                                  |
|--|--|--|--|
| # of Fixed counters per thread               | 3  | 3  | Use CPUID to enumerate # of counters.    |
| # of general-purpose counters per core       | 8  | 8  |  |
| Counter width (R,W)                          | R:48, W: 32/48   | R:48, W:32   | See Section 18.2.2.                      |
| # of programmable counters per thread        | 4 or (8 if a core not shared by two threads)   | 4  | Use CPUID to enumerate # of counters.    |
| PMI Overhead Mitigation                      | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_on_LBR with legacy semantics for branch profiling.</li> <li>▪ Freeze_while_SMM.</li> </ul> | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_on_LBR with legacy semantics for branch profiling.</li> <li>▪ Freeze_while_SMM.</li> </ul> | See Section 17.4.7.                      |
| Processor Event Based Sampling (PEBS) Events | See Table 18-12.   | See Table 18-68.   | IA32_PMC4-IA32_PMC7 do not support PEBS. |

**Table 18-10. Core PMU Comparison (Contd.)**

| Box                     | Intel® microarchitecture code name Sandy Bridge  | Intel® microarchitecture code name Westmere | Comment                     |
|-------------------------|--|---|-----------------------------|
| PEBS-Load Latency       | See Section 18.3.4.4.2;<br><ul style="list-style-type: none"> <li>▪ Data source encoding</li> <li>▪ STLB miss encoding</li> <li>▪ Lock transaction encoding</li> </ul> | Data source encoding                        |                             |
| PEBS-Precise Store      | Section 18.3.4.4.3   | No  |                             |
| PEBS-PDIR               | Yes (using precise INST_RETIRED.ALL).  | No  |                             |
| Off-core Response Event | MSR 1A6H and 1A7H, extended request and response types.  | MSR 1A6H and 1A7H, limited response types.  | Nehalem supports 1A6H only. |

**18.3.4.1 Global Counter Control Facilities In Intel® Microarchitecture Code Name Sandy Bridge**

The number of general-purpose performance counters visible to a logical processor can vary across Processors based on Intel microarchitecture code name Sandy Bridge. Software must use CPUID to determine the number performance counters/event select registers (See Section 18.2.1.1).



**Figure 18-26. IA32\_PERF\_GLOBAL\_CTRL MSR in Intel® Microarchitecture Code Name Sandy Bridge**

Figure 18-42 depicts the layout of IA32\_PERF\_GLOBAL\_CTRL MSR. The enable bits (PMC4\_EN, PMC5\_EN, PMC6\_EN, PMC7\_EN) corresponding to IA32\_PMC4-IA32\_PMC7 are valid only if CPUID.0AH:EAX[15:8] reports a value of '8'. If CPUID.0AH:EAX[15:8] = 4, attempts to set the invalid bits will cause #GP.

Each enable bit in IA32\_PERF\_GLOBAL\_CTRL is AND'ed with the enable bits for all privilege levels in the respective IA32\_PERFEVTSELx or IA32\_PERF\_FIXED\_CTR\_CTRL MSRs to start/stop the counting of respective counters. Counting is enabled if the AND'ed results is true; counting is disabled when the result is false.

IA32\_PERF\_GLOBAL\_STATUS MSR provides single-bit status used by software to query the overflow condition of each performance counter. IA32\_PERF\_GLOBAL\_STATUS[bit 62] indicates overflow conditions of the DS area data buffer (see Figure 18-27). A value of 1 in each bit of the PMCx\_OVF field indicates an overflow condition has occurred in the associated counter.

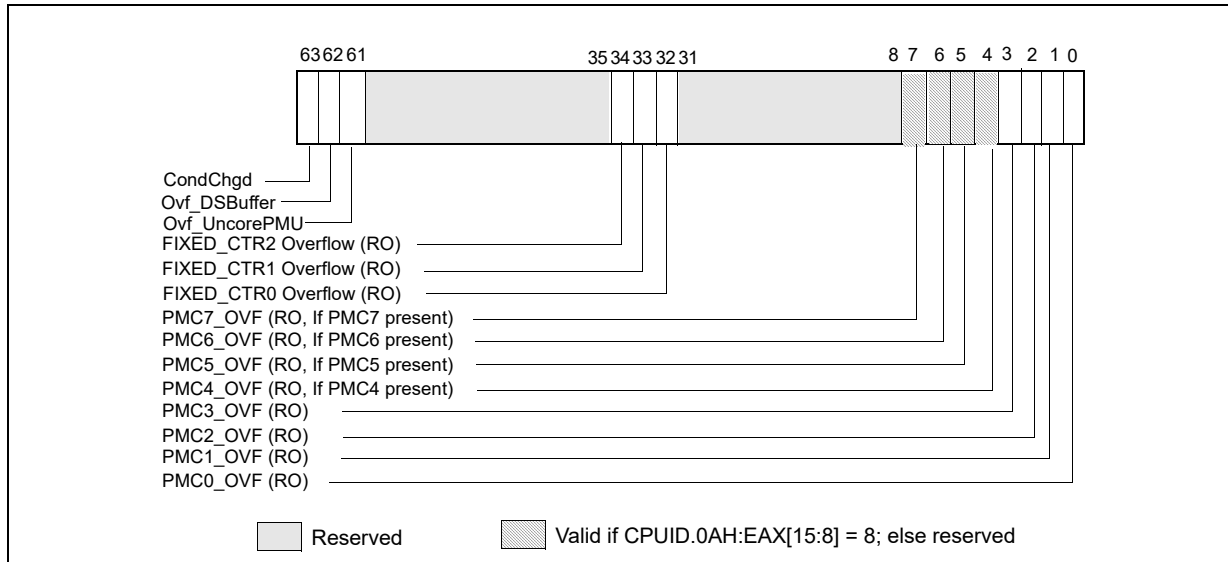


Figure 18-27. IA32\_PERF\_GLOBAL\_STATUS MSR in Intel® Microarchitecture Code Name Sandy Bridge

When a performance counter is configured for PEBS, an overflow condition in the counter will arm PEBS. On the subsequent event following overflow, the processor will generate a PEBS event. On a PEBS event, the processor will perform bounds checks based on the parameters defined in the DS Save Area (see Section 17.4.9). Upon successful bounds checks, the processor will store the data record in the defined buffer area, clear the counter overflow status, and reload the counter. If the bounds checks fail, the PEBS will be skipped entirely. In the event that the PEBS buffer fills up, the processor will set the OvfBuffer bit in MSR\_PERF\_GLOBAL\_STATUS.

IA32\_PERF\_GLOBAL\_OVF\_CTL MSR allows software to clear overflow the indicators for general-purpose or fixed-function counters via a single WRMSR (see Figure 18-28). Clear overflow indications when:

- Setting up new values in the event select and/or UMASK field for counting or interrupt based sampling.
- Reloading counter values to continue sampling.
- Disabling event counting or interrupt based sampling.




Figure 18-28. IA32\_PERF\_GLOBAL\_OVF\_CTRL MSR in Intel microarchitecture code name Sandy Bridge

### 18.3.4.2 Counter Coalescence

In processors based on Intel microarchitecture code name Sandy Bridge, each processor core implements eight general-purpose counters. CPUID.0AH:EAX[15:8] will report either 4 or 8 depending specific processor's product features.

If a processor core is shared by two logical processors, each logical processors can access 4 counters (IA32\_PMC0-IA32\_PMC3). This is the same as in the prior generation for processors based on Intel microarchitecture code name Nehalem.

If a processor core is not shared by two logical processors, all eight general-purpose counters are visible, and CPUID.0AH:EAX[15:8] reports 8. IA32\_PMC4-IA32\_PMC7 occupy MSR addresses 0C5H through 0C8H. Each counter is accompanied by an event select MSR (IA32\_PERFEVTSEL4-IA32\_PERFEVTSEL7).

If CPUID.0AH:EAX[15:8] report 4, access to IA32\_PMC4-IA32\_PMC7, IA32\_PMC4-IA32\_PMC7 will cause #GP. Writing 1's to bit position 7:4 of IA32\_PERF\_GLOBAL\_CTRL, IA32\_PERF\_GLOBAL\_STATUS, or IA32\_PERF\_GLOBAL\_OVF\_CTL will also cause #GP.

### 18.3.4.3 Full Width Writes to Performance Counters

Processors based on Intel microarchitecture code name Sandy Bridge support full-width writes to the general-purpose counters, IA32\_PMCx. Support of full-width writes are enumerated by IA32\_PERF\_CAPABILITIES.FW\_WRITES[13] (see Section 18.2.4).

The default behavior of IA32\_PMCx is unchanged, i.e. WRMSR to IA32\_PMCx results in a sign-extended 32-bit value of the input EAX written into IA32\_PMCx. Full-width writes must issue WRMSR to a dedicated alias MSR address for each IA32\_PMCx.

Software must check the presence of full-width write capability and the presence of the alias address IA32\_A\_PMCx by testing IA32\_PERF\_CAPABILITIES[13].

### 18.3.4.4 PEBS Support in Intel® Microarchitecture Code Name Sandy Bridge

Processors based on Intel microarchitecture code name Sandy Bridge support PEBS, similar to those offered in prior generation, with several enhanced features. The key components and differences of PEBS facility relative to Intel microarchitecture code name Westmere is summarized in Table 18-11.

**Table 18-11. PEBS Facility Comparison**

| Box                       | Intel® microarchitecture code name Sandy Bridge                         | Intel® microarchitecture code name Westmere | Comment                                   |
|---------------------------|---|---|---|
| Valid IA32_PMCx           | PMCO-PMC3   | PMCO-PMC3                                   | No PEBS on PMC4-PMC7.                     |
| PEBS Buffer Programming   | Section 18.3.1.1.1  | Section 18.3.1.1.1                          | Unchanged                                 |
| IA32_PEBS_ENABLE Layout   | Figure 18-29  | Figure 18-15                                |   |
| PEBS record layout        | Physical Layout same as Table 18-3.                                     | Table 18-3                                  | Enhanced fields at offsets 98H, A0H, A8H. |
| PEBS Events               | See Table 18-12.  | See Table 18-68.                            | IA32_PMC4-IA32_PMC7 do not support PEBS.  |
| PEBS-Load Latency         | See Table 18-13.  | Table 18-4                                  |   |
| PEBS-Precise Store        | Yes; see Section 18.3.4.4.3.  | No  | IA32_PMC3 only                            |
| PEBS-PDIR                 | Yes   | No  | IA32_PMC1 only                            |
| PEBS skid from EventingIP | 1 (or 2 if micro+macro fusion)  | 1   |   |
| SAMPLING Restriction      | Small SAV(CountDown) value incur higher overhead than prior generation. |   |   |

Only IA32\_PMC0 through IA32\_PMC3 support PEBS.

### NOTE

PEBS events are only valid when the following fields of IA32\_PERFEVTSELx are all zero: AnyThread, Edge, Invert, CMask.

In a PMU with PDIR capability, PEBS behavior is unpredictable if IA32\_PERFEVTSELx or IA32\_PMCx is changed for a PEBS-enabled counter while an event is being counted. To avoid this, changes to the programming or value of a PEBS-enabled counter should be performed when the counter is disabled.

In IA32\_PEBS\_ENABLE MSR, bit 63 is defined as PS\_ENABLE: When set, this enables IA32\_PMC3 to capture precise store information. Only IA32\_PMC3 supports the precise store facility. In typical usage of PEBS, the bit fields in IA32\_PEBS\_ENABLE are written to when the agent software starts PEBS operation; the enabled bit fields should be modified only when re-programming another PEBS event or cleared when the agent uses the performance counters for non-PEBS operations.

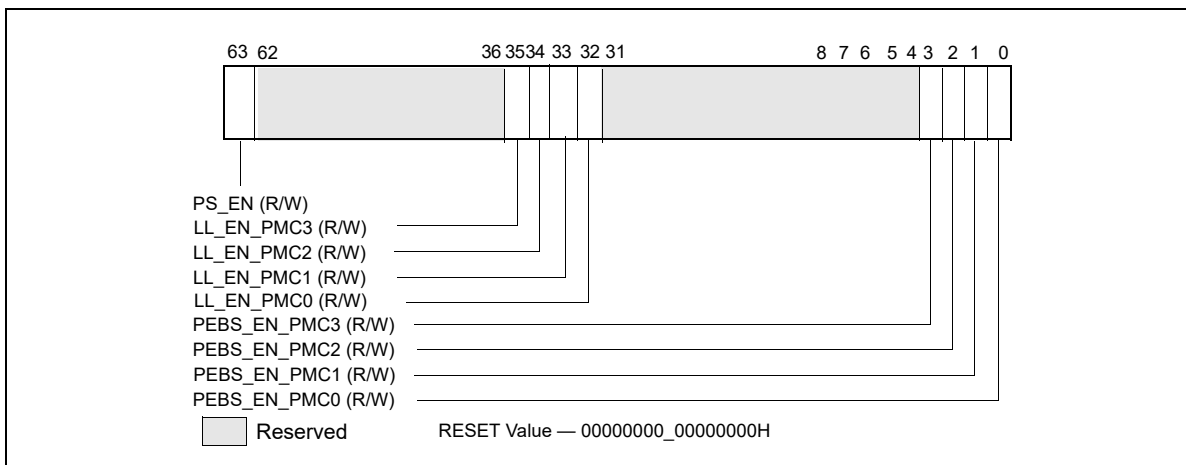


Figure 18-29. Layout of IA32\_PEBS\_ENABLE MSR

#### 18.3.4.4.1 PEBS Record Format

The layout of PEBS records physically identical to those shown in Table 18-3, but the fields at offset 98H, A0H and A8H have been enhanced to support additional PEBS capabilities.

- Load/Store Data Linear Address (Offset 98H): This field will contain the linear address of the source of the load, or linear address of the destination of the store.
- Data Source /Store Status (Offset A0H): When load latency is enabled, this field will contain three piece of information (including an encoded value indicating the source which satisfied the load operation). The source field encodings are detailed in Table 18-4. When precise store is enabled, this field will contain information indicating the status of the store, as detailed in Table 19.
- Latency Value/0 (Offset A8H): When load latency is enabled, this field contains the latency in cycles to service the load. This field is not meaningful when precise store is enabled and will be written to zero in that case. Upon writing the PEBS record, microcode clears the overflow status bits in the IA32\_PERF\_GLOBAL\_STATUS corresponding to those counters that both overflowed and were enabled in the IA32\_PEBS\_ENABLE register. The status bits of other counters remain unaffected.

The number PEBS events has expanded. The list of PEBS events supported in Intel microarchitecture code name Sandy Bridge is shown in Table 18-12.

**Table 18-12. PEBS Performance Events for Intel® Microarchitecture Code Name Sandy Bridge**

| Event Name                    | Event Select | Sub-event       | UMask            |
|-------------------------------|--------------|-----------------|------------------|
| INST_RETIRED                  | C0H          | PREC_DIST       | 01H <sup>1</sup> |
| UOPS_RETIRED                  | C2H          | All             | 01H              |
|                               |              | Retire_Slots    | 02H              |
| BR_INST_RETIRED               | C4H          | Conditional     | 01H              |
|                               |              | Near_Call       | 02H              |
|                               |              | All_branches    | 04H              |
|                               |              | Near_Return     | 08H              |
|                               |              | Near_Taken      | 20H              |
| BR_MISP_RETIRED               | C5H          | Conditional     | 01H              |
|                               |              | Near_Call       | 02H              |
|                               |              | All_branches    | 04H              |
|                               |              | Not_Taken       | 10H              |
|                               |              | Taken           | 20H              |
| MEM_UOPS_RETIRED              | D0H          | STLB_MISS_LOADS | 11H              |
|                               |              | STLB_MISS_STORE | 12H              |
|                               |              | LOCK_LOADS      | 21H              |
|                               |              | SPLIT_LOADS     | 41H              |
|                               |              | SPLIT_STORES    | 42H              |
|                               |              | ALL_LOADS       | 81H              |
|                               |              | ALL_STORES      | 82H              |
| MEM_LOAD_UOPS_RETIRED         | D1H          | L1_Hit          | 01H              |
|                               |              | L2_Hit          | 02H              |
|                               |              | L3_Hit          | 04H              |
|                               |              | Hit_LFB         | 40H              |
| MEM_LOAD_UOPS_LLC_HIT_RETIRED | D2H          | XSNP_Miss       | 01H              |
|                               |              | XSNP_Hit        | 02H              |
|                               |              | XSNP_Hitm       | 04H              |
|                               |              | XSNP_None       | 08H              |

**NOTES:**

1. Only available on IA32\_PMC1.

**18.3.4.4.2 Load Latency Performance Monitoring Facility**

The load latency facility in Intel microarchitecture code name Sandy Bridge is similar to that in prior microarchitecture. It provides software a means to characterize the average load latency to different levels of cache/memory hierarchy. This facility requires processor supporting enhanced PEBS record format in the PEBS buffer, see Table 18-3 and Section 18.3.4.4.1. This field measures the load latency from load's first dispatch of till final data writeback from the memory subsystem. The latency is reported for retired demand load operations and in core cycles (it accounts for re-dispatches).

To use this feature software must assure:

- One of the IA32\_PERFEVTSELx MSR is programmed to specify the event unit MEM\_TRANS\_RETIRED, and the LATENCY\_ABOVE\_THRESHOLD event mask must be specified (IA32\_PerfEvtSelX[15:0] = 1CDH). The corresponding counter IA32\_PMCx will accumulate event counts for architecturally visible loads which exceed the programmed latency threshold specified separately in a MSR. Stores are ignored when this event is

programmed. The CMASK or INV fields of the IA32\_PerfEvtSelX register used for counting load latency must be 0. Writing other values will result in undefined behavior.

- The MSR\_PEBS\_LD\_LAT\_THRESHOLD MSR is programmed with the desired latency threshold in core clock cycles. Loads with latencies greater than this value are eligible for counting and latency data reporting. The minimum value that may be programmed in this register is 3 (the minimum detectable load latency is 4 core clock cycles).
- The PEBS enable bit in the IA32\_PEBS\_ENABLE register is set for the corresponding IA32\_PMCx counter register. This means that both the PEBS\_EN\_CTRX and LL\_EN\_CTRX bits must be set for the counter(s) of interest. For example, to enable load latency on counter IA32\_PMC0, the IA32\_PEBS\_ENABLE register must be programmed with the 64-bit value 00000001.00000001H.
- When Load latency event is enabled, no other PEBS event can be configured with other counters.

When the load-latency facility is enabled, load operations are randomly selected by hardware and tagged to carry information related to data source locality and latency. Latency and data source information of tagged loads are updated internally. The MEM\_TRANS\_RETIRE event for load latency counts only tagged retired loads. If a load is cancelled it will not be counted and the internal state of the load latency facility will not be updated. In this case the hardware will tag the next available load.

When a PEBS assist occurs, the last update of latency and data source information are captured by the assist and written as part of the PEBS record. The PEBS sample after value (SAV), specified in PEBS CounterX Reset, operates orthogonally to the tagging mechanism. Loads are randomly tagged to collect latency data. The SAV controls the number of tagged loads with latency information that will be written into the PEBS record field by the PEBS assists. The load latency data written to the PEBS record will be for the last tagged load operation which retired just before the PEBS assist was invoked.

The physical layout of the PEBS records is the same as shown in Table 18-3. The specificity of Data Source entry at offset A0H has been enhanced to report three pieces of information.

**Table 18-13. Layout of Data Source Field of Load Latency Record**

| Field     | Position | Description  |
|-----------|----------|--|
| Source    | 3:0      | See Table 18-4   |
| STLB_MISS | 4        | 0: The load did not miss the STLB (hit the DTLB or STLB).<br>1: The load missed the STLB.          |
| Lock      | 5        | 0: The load was not part of a locked transaction.<br>1: The load was part of a locked transaction. |
| Reserved  | 63:6     | Reserved   |

The layout of MSR\_PEBS\_LD\_LAT\_THRESHOLD is the same as shown in Figure 18-17.

#### 18.3.4.4.3 Precise Store Facility

Processors based on Intel microarchitecture code name Sandy Bridge offer a precise store capability that complements the load latency facility. It provides a means to profile store memory references in the system.

Precise stores leverage the PEBS facility and provide additional information about sampled stores. Having precise memory reference events with linear address information for both loads and stores can help programmers improve data structure layout, eliminate remote node references, and identify cache-line conflicts in NUMA systems.

Only IA32\_PMC3 can be used to capture precise store information. After enabling this facility, counter overflows will initiate the generation of PEBS records as previously described in PEBS. Upon counter overflow hardware captures the linear address and other status information of the next store that retires. This information is then written to the PEBS record.

To enable the precise store facility, software must complete the following steps. Please note that the precise store facility relies on the PEBS facility, so the PEBS configuration requirements must be completed before attempting to capture precise store information.

- Complete the PEBS configuration steps.

- Program the MEM\_TRANS\_RETIREDPRECISE\_STORE event in IA32\_PERFVTSEL3. Only counter 3 (IA32\_PMC3) supports collection of precise store information.
- Set IA32\_PEBS\_ENABLE[3] and IA32\_PEBS\_ENABLE[63]. This enables IA32\_PMC3 as a PEBS counter and enables the precise store facility, respectively.

The precise store information written into a PEBS record affects entries at offset 98H, A0H and A8H of Table 18-3. The specificity of Data Source entry at offset A0H has been enhanced to report three piece of information.

**Table 18-14. Layout of Precise Store Information In PEBS Record**

| Field                     | Offset | Description   |
|---------------------------|--------|---|
| Store Data Linear Address | 98H    | The linear address of the destination of the store.   |
| Store Status              | A0H    | <p><b>L1D Hit</b> (Bit 0): The store hit the data cache closest to the core (lowest latency cache) if this bit is set, otherwise the store missed the data cache.</p> <p><b>STLB Miss</b> (bit 4): The store missed the STLB if set, otherwise the store hit the STLB</p> <p><b>Locked Access</b> (bit 5): The store was part of a locked access if set, otherwise the store was not part of a locked access.</p> |
| Reserved                  | A8H    | Reserved  |

#### 18.3.4.4.4 Precise Distribution of Instructions Retired (PDIR)

Upon triggering a PEBS assist, there will be a finite delay between the time the counter overflows and when the microcode starts to carry out its data collection obligations. INST\_RETIREDP is a very common event that is used to sample where performance bottleneck happened and to help identify its location in instruction address space. Even if the delay is constant in core clock space, it invariably manifest as variable “skids” in instruction address space. This creates a challenge for programmers to profile a workload and pinpoint the location of bottlenecks.

The core PMU in processors based on Intel microarchitecture code name Sandy Bridge include a facility referred to as precise distribution of Instruction Retired (PDIR).

The PDIR facility mitigates the “skid” problem by providing an early indication of when the INST\_RETIREDP counter is about to overflow, allowing the machine to more precisely trap on the instruction that actually caused the counter overflow thus eliminating skid.

PDIR applies only to the INST\_RETIREDP.ALL precise event, and must use IA32\_PMC1 with PerfEvtSel1 property configured and bit 1 in the IA32\_PEBS\_ENABLE set to 1. INST\_RETIREDP.ALL is a non-architectural performance event, it is not supported in prior generation microarchitectures. Additionally, on processors with CPUID DisplayFamily\_DisplayModel signatures of 06\_2A and 06\_2D, the tool that programs PDIR should quiesce the rest of the programmable counters in the core when PDIR is active.

#### 18.3.4.5 Off-core Response Performance Monitoring


The core PMU in processors based on Intel microarchitecture code name Sandy Bridge provides off-core response facility similar to prior generation. Off-core response can be programmed only with a specific pair of event select and counter MSR, and with specific event codes and predefine mask bit value in a dedicated MSR to specify attributes of the off-core transaction. Two event codes are dedicated for off-core response event programming. Each event code for off-core response monitoring requires programming an associated configuration MSR, MSR\_OFFCORE\_RSP\_x. Table 18-15 lists the event code, mask value and additional off-core configuration MSR that must be programmed to count off-core response events using IA32\_PMCx.



**Table 18-15. Off-Core Response Event Encoding**

| Counter | Event code | UMask | Required Off-core Response MSR   |
|---------|------------|-------|----------------------------------|
| PMCO-3  | B7H        | 01H   | MSR_OFFCORE_RSP_0 (address 1A6H) |
| PMCO-3  | BBH        | 01H   | MSR_OFFCORE_RSP_1 (address 1A7H) |

The layout of MSR\_OFFCORE\_RSP\_0 and MSR\_OFFCORE\_RSP\_1 are shown in Figure 18-30 and Figure 18-31. Bits 15:0 specifies the request type of a transaction request to the uncore. Bits 30:16 specifies supplier information, bits 37:31 specifies snoop response information.



**Figure 18-30. Request\_Type Fields for MSR\_OFFCORE\_RSP\_x**

**Table 18-16. MSR\_OFFCORE\_RSP\_x Request\_Type Field Definition**

| Bit Name       | Offset | Description  |
|----------------|--------|--|
| DMND_DATA_RD   | 0      | (R/W). Counts the number of demand data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count L2 data read prefetches or instruction fetches. |
| DMND_RFO       | 1      | (R/W). Counts the number of demand and DCU prefetch reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO prefetches.                                     |
| DMND_IFETCH    | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.   |
| WB             | 3      | (R/W). Counts the number of writeback (modified to exclusive) transactions.  |
| PF_DATA_RD     | 4      | (R/W). Counts the number of data cacheline reads generated by L2 prefetchers.  |
| PF_RFO         | 5      | (R/W). Counts the number of RFO requests generated by L2 prefetchers.  |
| PF_IFETCH      | 6      | (R/W). Counts the number of code reads generated by L2 prefetchers.  |
| PF_LLC_DATA_RD | 7      | (R/W). L2 prefetcher to L3 for loads.  |
| PF_LLC_RFO     | 8      | (R/W). RFO requests generated by L2 prefetcher   |
| PF_LLC_IFETCH  | 9      | (R/W). L2 prefetcher to L3 for instruction fetches.  |
| BUS_LOCKS      | 10     | (R/W). Bus lock and split lock requests  |
| STRM_ST        | 11     | (R/W). Streaming store requests  |
| OTHER          | 15     | (R/W). Any other request that crosses IDI, including I/O.  |

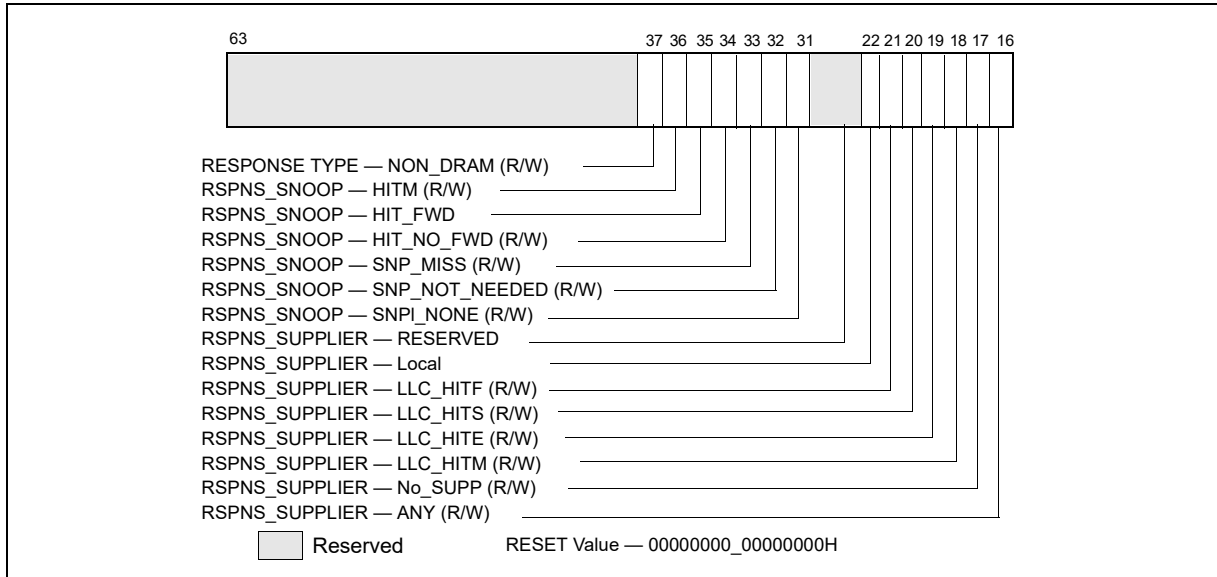


Figure 18-31. Response\_Supplier and Snoop Info Fields for MSR\_OFFCORE\_RSP\_x

To properly program this extra register, software must set at least one request type bit and a valid response type pattern. Otherwise, the event count reported will be zero. It is permissible and useful to set multiple request and response type bits in order to obtain various classes of off-core response events. Although MSR\_OFFCORE\_RSP\_x allow an agent software to program numerous combinations that meet the above guideline, not all combinations produce meaningful data.

Table 18-17. MSR\_OFFCORE\_RSP\_x Response Supplier Info Field Definition

| Subtype       | Bit Name | Offset | Description                                    |
|---------------|----------|--------|--|
| Common        | Any      | 16     | (R/W). Catch all value for any response types. |
| Supplier Info | NO_SUPP  | 17     | (R/W). No Supplier Information available       |
|               | LLC_HITM | 18     | (R/W). M-state initial lookup stat in L3.      |
|               | LLC_HITE | 19     | (R/W). E-state                                 |
|               | LLC_HITS | 20     | (R/W). S-state                                 |
|               | LLC_HITF | 21     | (R/W). F-state                                 |
|               | LOCAL    | 22     | (R/W). Local DRAM Controller                   |
|               | Reserved |        | 30:23  |

To specify a complete offcore response filter, software must properly program bits in the request and response type fields. A valid request type must have at least one bit set in the non-reserved bits of 15:0. A valid response type must be a non-zero value of the following expression:

ANY | [(‘OR’ of Supplier Info Bits) & (‘OR’ of Snoop Info Bits)]

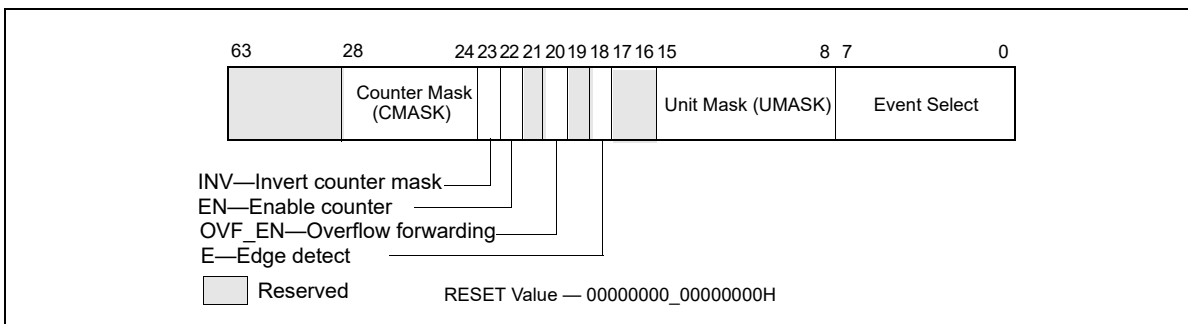
If “ANY” bit is set, the supplier and snoop info bits are ignored.

**Table 18-18. MSR\_OFFCORE\_RSP\_x Snoop Info Field Definition**

| Subtype    | Bit Name       | Offset | Description  |
|------------|----------------|--------|--|
| Snoop Info | SNP_NONE       | 31     | (R/W). No details on snoop-related information   |
|            | SNP_NOT_NEEDED | 32     | (R/W). No snoop was needed to satisfy the request.   |
|            | SNP_MISS       | 33     | (R/W). A snoop was needed and it missed all snooped caches:<br>-For LLC Hit, ReslHitl was returned by all cores<br>-For LLC Miss, Rspl was returned by all sockets and data was returned from DRAM.  |
|            | SNP_NO_FWD     | 34     | (R/W). A snoop was needed and it hits in at least one snooped cache. Hit denotes a cache-line was valid before snoop effect. This includes:<br>-Snoop Hit w/ Invalidation (LLC Hit, RFO)<br>-Snoop Hit, Left Shared (LLC Hit/Miss, IFetch/Data_RD)<br>-Snoop Hit w/ Invalidation and No Forward (LLC Miss, RFO Hit S)<br>In the LLC Miss case, data is returned from DRAM. |
|            | SNP_FWD        | 35     | (R/W). A snoop was needed and data was forwarded from a remote socket. This includes:<br>-Snoop Forward Clean, Left Shared (LLC Hit/Miss, IFetch/Data_RD/RFT).   |
|            | HITM           | 36     | (R/W). A snoop was needed and it HitM-ed in local or remote cache. HitM denotes a cache-line was in modified state before effect as a results of snoop. This includes:<br>-Snoop HitM w/ WB (LLC miss, IFetch/Data_RD)<br>-Snoop Forward Modified w/ Invalidation (LLC Hit/Miss, RFO)<br>-Snoop MtoS (LLC Hit, IFetch/Data_RD).  |
|            | NON_DRAM       | 37     | (R/W). Target was non-DRAM system address. This includes MMIO transactions.  |

**18.3.4.6 Uncore Performance Monitoring Facilities In Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series**

The uncore sub-system in Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series provides a unified L3 that can support up to four processor cores. The L3 cache consists multiple slices, each slice interface with a processor via a coherence engine, referred to as a C-Box. Each C-Box provides dedicated facility of MSRs to select uncore performance monitoring events and each C-Box event select MSR is paired with a counter register, similar in style as those described in Section 18.3.1.2.2. The ARB unit in the uncore also provides its local performance counters and event select MSRs. The layout of the event select MSRs in the C-Boxes and the ARB unit are shown in Figure 18-32.



**Figure 18-32. Layout of Uncore PERFVTSSEL MSR for a C-Box Unit or the ARB Unit**

The bit fields of the uncore event select MSRs for a C-box unit or the ARB unit are summarized below:

- Event\_Select (bits 7:0) and UMASK (bits 15:8): Specifies the microarchitectural condition to count in a local uncore PMU counter, see Table 19-18.
- E (bit 18): Enables edge detection filtering, if 1.
- OVF\_EN (bit 20): Enables the overflow indicator from the uncore counter forwarded to MSR\_UNC\_PERF\_GLOBAL\_CTRL, if 1.
- EN (bit 22): Enables the local counter associated with this event select MSR.
- INV (bit 23): Event count increments with non-negative value if 0, with negated value if 1.
- CMASK (bits 28:24): Specifies a positive threshold value to filter raw event count input.

At the uncore domain level, there is a master set of control MSRs that centrally manages all the performance monitoring facility of uncore units. Figure 18-33 shows the layout of the uncore domain global control.

When an uncore counter overflows, a PMI can be routed to a processor core. Bits 3:0 of MSR\_UNC\_PERF\_GLOBAL\_CTRL can be used to select which processor core to handle the uncore PMI. Software must then write to bit 13 of IA32\_DEBUGCTL (at address 1D9H) to enable this capability.

- PMI\_SEL\_Core#: Enables the forwarding of an uncore PMI request to a processor core, if 1. If bit 30 (WakePMI) is '1', a wake request is sent to the respective processor core prior to sending the PMI.
- EN: Enables the fixed uncore counter, the ARB counters, and the CBO counters in the uncore PMU, if 1. This bit is cleared if bit 31 (FREEZE) is set and any enabled uncore counters overflow.
- WakePMI: Controls sending a wake request to any halted processor core before issuing the uncore PMI request. If a processor core was halted and not sent a wake request, the uncore PMI will not be serviced by the processor core.
- FREEZE: Provides the capability to freeze all uncore counters when an overflow condition occurs in a unit counter. When this bit is set, and a counter overflow occurs, the uncore PMU logic will clear the global enable bit (bit 29).




Figure 18-33. Layout of MSR\_UNC\_PERF\_GLOBAL\_CTRL MSR for Uncore

Additionally, there is also a fixed counter, counting uncore clockticks, for the uncore domain. Table 18-19 summarizes the number MSR for uncore PMU for each box.

**Table 18-19. Uncore PMU MSR Summary**

| Box           | # of Boxes   | Counters per Box | Counter Width | General Purpose | Global Enable | Comment                                       |
|---------------|--------------|------------------|---------------|-----------------|---------------|---|
| C-Box         | SKU specific | 2                | 44            | Yes             | Per-box       | Up to 4, see Table 2-20<br>MSR_UNC_CBO_CONFIG |
| ARB           | 1            | 2                | 44            | Yes             | Uncore        |   |
| Fixed Counter | N.A.         | N.A.             | 48            | No              | Uncore        |   |

#### 18.3.4.6.1 Uncore Performance Monitoring Events

There are certain restrictions on the uncore performance counters in each C-Box. Specifically,

- Occupancy events are supported only with counter 0 but not counter 1.

Other uncore C-Box events can be programmed with either counter 0 or 1.

The C-Box uncore performance events described in Table 19-18 can collect performance characteristics of transactions initiated by processor core. In that respect, they are similar to various sub-events in the OFFCORE\_RESPONSE family of performance events in the core PMU. Information such as data supplier locality (LLC HIT/MISS) and snoop responses can be collected via OFFCORE\_RESPONSE and qualified on a per-thread basis.

On the other hand, uncore performance event logic can not associate its counts with the same level of per-thread qualification attributes as the core PMU events can. Therefore, whenever similar event programming capabilities are available from both core PMU and uncore PMU, the recommendation is that utilizing the core PMU events may be less affected by artifacts, complex interactions and other factors.

#### 18.3.4.7 Intel® Xeon® Processor E5 Family Performance Monitoring Facility

The Intel® Xeon® Processor E5 Family (and Intel® Core™ i7-3930K Processor) are based on Intel microarchitecture code name Sandy Bridge-E. While the processor cores share the same microarchitecture as those of the Intel® Xeon® Processor E3 Family and 2nd generation Intel Core i7-2xxx, Intel Core i5-2xxx, Intel Core i3-2xxx processor series, the uncore subsystems are different. An overview of the uncore performance monitoring facilities of the Intel Xeon processor E5 family (and Intel Core i7-3930K processor) is described in Section 18.3.4.8.

Thus, the performance monitoring facilities in the processor core generally are the same as those described in Section 18.6.3 through Section 18.3.4.5. However, the MSR\_OFFCORE\_RSP\_0/MSR\_OFFCORE\_RSP\_1 Response Supplier Info field shown in Table 18-17 applies to Intel Core Processors with CPUID signature of DisplayFamily\_DisplayModel encoding of 06\_2AH; Intel Xeon processor with CPUID signature of DisplayFamily\_DisplayModel encoding of 06\_2DH supports an additional field for remote DRAM controller shown in Table 18-20. Additionally, there are some small differences in the non-architectural performance monitoring events (see Table 19-16).

**Table 18-20. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definitions**

| Subtype       | Bit Name | Offset | Description   |
|---------------|----------|--------|---|
| Common        | Any      | 16     | (R/W). Catch all value for any response types.          |
| Supplier Info | NO_SUPP  | 17     | (R/W). No Supplier Information available                |
|               | LLC_HITM | 18     | (R/W). M-state initial lookup stat in L3.               |
|               | LLC_HITE | 19     | (R/W). E-state  |
|               | LLC_HITS | 20     | (R/W). S-state  |
|               | LLC_HITF | 21     | (R/W). F-state  |
|               | LOCAL    | 22     | (R/W). Local DRAM Controller                            |
|               | Remote   | 30:23  | (R/W): Remote DRAM Controller (either all 0s or all 1s) |

### 18.3.4.8 Intel® Xeon® Processor E5 Family Uncore Performance Monitoring Facility

The uncore subsystem in the Intel Xeon processor E5-2600 product family has some similarities with those of the Intel Xeon processor E7 family. Within the uncore subsystem, localized performance counter sets are provided at logic control unit scope. For example, each Cbox caching agent has a set of local performance counters, and the power controller unit (PCU) has its own local performance counters. Up to 8 C-Box units are supported in the uncore sub-system.

Table 18-21 summarizes the uncore PMU facilities providing MSR interfaces.

**Table 18-21. Uncore PMU MSR Summary for Intel® Xeon® Processor E5 Family**

| Box   | # of Boxes | Counters per Box | Counter Width | General Purpose | Global Enable | Sub-control MSRs |
|-------|------------|------------------|---------------|-----------------|---------------|------------------|
| C-Box | 8          | 4                | 44            | Yes             | per-box       | None             |
| PCU   | 1          | 4                | 48            | Yes             | per-box       | Match/Mask       |
| U-Box | 1          | 2                | 44            | Yes             | uncore        | None             |

Details of the uncore performance monitoring facility of Intel Xeon Processor E5 family is available in "Intel® Xeon® Processor E5 Uncore Performance Monitoring Programming Reference Manual". The MSR-based uncore PMU interfaces are listed in Table 2-23.

### 18.3.5 3rd Generation Intel® Core™ Processor Performance Monitoring Facility

The 3rd generation Intel® Core™ processor family and Intel® Xeon® processor E3-1200v2 product family are based on the Ivy Bridge microarchitecture. The performance monitoring facilities in the processor core generally are the same as those described in Section 18.6.3 through Section 18.3.4.5. The non-architectural performance monitoring events supported by the processor core are listed in Table 19-16.

#### 18.3.5.1 Intel® Xeon® Processor E5 v2 and E7 v2 Family Uncore Performance Monitoring Facility

The uncore subsystem in the Intel Xeon processor E5 v2 and Intel Xeon Processor E7 v2 product families are based on the Ivy Bridge-E microarchitecture. There are some similarities with those of the Intel Xeon processor E5 family based on the Sandy Bridge microarchitecture. Within the uncore subsystem, localized performance counter sets are provided at logic control unit scope.

Details of the uncore performance monitoring facility of Intel Xeon Processor E5 v2 and Intel Xeon Processor E7 v2 families are available in "Intel® Xeon® Processor E5 v2 and E7 v2 Uncore Performance Monitoring Programming Reference Manual". The MSR-based uncore PMU interfaces are listed in Table 2-27.

### 18.3.6 4th Generation Intel® Core™ Processor Performance Monitoring Facility

The 4th generation Intel® Core™ processor and Intel® Xeon® processor E3-1200 v3 product family are based on the Haswell microarchitecture. The core PMU supports architectural performance monitoring capability with version ID 3 (see Section 18.2.3) and a host of non-architectural monitoring capabilities.

Architectural performance monitoring version 3 capabilities are described in Section 18.2.3.

The core PMU's capability is similar to those described in Section 18.6.3 through Section 18.3.4.5, with some differences and enhancements summarized in Table 18-22. Additionally, the core PMU provides some enhancement to support performance monitoring when the target workload contains instruction streams using Intel® Transactional Synchronization Extensions (TSX), see Section 18.3.6.5. For details of Intel TSX, see Chapter 16, "Programming with Intel® Transactional Synchronization Extensions" of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*.

**Table 18-22. Core PMU Comparison**

| Box  | Intel® microarchitecture code name Haswell   | Intel® microarchitecture code name Sandy Bridge  | Comment                                  |
|--|--|--|--|
| # of Fixed counters per thread               | 3  | 3  |  |
| # of general-purpose counters per core       | 8  | 8  |  |
| Counter width (R,W)                          | R:48, W: 32/48   | R:48, W: 32/48   | See Section 18.2.2.                      |
| # of programmable counters per thread        | 4 or (8 if a core not shared by two threads)   | 4 or (8 if a core not shared by two threads)   | Use CPUID to enumerate # of counters.    |
| PMI Overhead Mitigation                      | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_on_LBR with legacy semantics for branch profiling.</li> <li>▪ Freeze_while_SMM.</li> </ul> | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_on_LBR with legacy semantics for branch profiling.</li> <li>▪ Freeze_while_SMM.</li> </ul> | See Section 17.4.7.                      |
| Processor Event Based Sampling (PEBS) Events | See Table 18-12 and Section 18.3.6.5.1.  | See Table 18-12.   | IA32_PMC4-IA32_PMC7 do not support PEBS. |
| PEBS-Load Latency                            | See Section 18.3.4.4.2.  | See Section 18.3.4.4.2.  |  |
| PEBS-Precise Store                           | No, replaced by Data Address profiling.  | Section 18.3.4.4.3   |  |
| PEBS-PDIR                                    | Yes (using precise INST_RETIRED.ALL)   | Yes (using precise INST_RETIRED.ALL)   |  |
| PEBS-EventingIP                              | Yes  | No   |  |
| Data Address Profiling                       | Yes  | No   |  |
| LBR Profiling                                | Yes  | Yes  |  |
| Call Stack Profiling                         | Yes, see Section 17.11.  | No   | Use LBR facility.                        |
| Off-core Response Event                      | MSR 1A6H and 1A7H; extended request and response types.  | MSR 1A6H and 1A7H; extended request and response types.  |  |
| Intel TSX support for Perfmon                | See Section 18.3.6.5.  | No   |  |

#### 18.3.6.1 Processor Event Based Sampling (PEBS) Facility

The PEBS facility in the 4th Generation Intel Core processor is similar to those in processors based on Intel microarchitecture code name Sandy Bridge, with several enhanced features. The key components and differences of PEBS facility relative to Intel microarchitecture code name Sandy Bridge is summarized in Table 18-23.

**Table 18-23. PEBS Facility Comparison**

| Box                       | Intel® microarchitecture code name Haswell                              | Intel® microarchitecture code name Sandy Bridge       | Comment                                  |
|---------------------------|---|---|--|
| Valid IA32_PMCx           | PMC0-PMC3   | PMC0-PMC3   | No PEBS on PMC4-PMC7                     |
| PEBS Buffer Programming   | Section 18.3.1.1.1  | Section 18.3.1.1.1                                    | Unchanged                                |
| IA32_PEBS_ENABLE Layout   | Figure 18-15  | Figure 18-29  |  |
| PEBS record layout        | Table 18-24; enhanced fields at offsets 98H, A0H, A8H, B0H.             | Table 18-3; enhanced fields at offsets 98H, A0H, A8H. |  |
| Precise Events            | See Table 18-12.  | See Table 18-12.                                      | IA32_PMC4-IA32_PMC7 do not support PEBS. |
| PEBS-Load Latency         | See Table 18-13.  | Table 18-13   |  |
| PEBS-Precise Store        | No, replaced by data address profiling.                                 | Yes; see Section 18.3.4.4.3.                          |  |
| PEBS-PDIR                 | Yes   | Yes   | IA32_PMC1 only.                          |
| PEBS skid from EventingIP | 1 (or 2 if micro+macro fusion)  | 1   |  |
| SAMPLING Restriction      | Small SAV(CountDown) value incur higher overhead than prior generation. |   |  |

Only IA32\_PMC0 through IA32\_PMC3 support PEBS.

#### NOTE

PEBS events are only valid when the following fields of IA32\_PERFEVTSELx are all zero: AnyThread, Edge, Invert, CMask.

In a PMU with PDIR capability, PEBS behavior is unpredictable if IA32\_PERFEVTSELx or IA32\_PMCx is changed for a PEBS-enabled counter while an event is being counted. To avoid this, changes to the programming or value of a PEBS-enabled counter should be performed when the counter is disabled.

### 18.3.6.2 PEBS Data Format

The PEBS record format for the 4th Generation Intel Core processor is shown in Table 18-24. The PEBS record format, along with debug/store area storage format, does not change regardless of whether IA-32e mode is active or not. CPUID.01H:ECX.DTES64[bit 2] reports whether the processor's DS storage format support is mode-independent. When set, it uses 64-bit DS storage format.



**Table 18-24. PEBS Record Format for 4th Generation Intel Core Processor Family**

| Byte Offset | Field    | Byte Offset | Field                                     |
|-------------|----------|-------------|---|
| 00H         | R/EFLAGS | 60H         | R10                                       |
| 08H         | R/EIP    | 68H         | R11                                       |
| 10H         | R/EAX    | 70H         | R12                                       |
| 18H         | R/EBX    | 78H         | R13                                       |
| 20H         | R/ECX    | 80H         | R14                                       |
| 28H         | R/EDX    | 88H         | R15                                       |
| 30H         | R/ESI    | 90H         | IA32_PERF_GLOBAL_STATUS                   |
| 38H         | R/EDI    | 98H         | Data Linear Address                       |
| 40H         | R/EBP    | A0H         | Data Source Encoding                      |
| 48H         | R/ESP    | A8H         | Latency value (core cycles)               |
| 50H         | R8       | B0H         | EventingIP                                |
| 58H         | R9       | B8H         | TX Abort Information (Section 18.3.6.5.1) |

The layout of PEBS records are almost identical to those shown in Table 18-3. Offset B0H is a new field that records the eventing IP address of the retired instruction that triggered the PEBS assist.

The PEBS records at offsets 98H, A0H, and ABH record data gathered from three of the PEBS capabilities in prior processor generations: load latency facility (Section 18.3.4.4.2), PDIR (Section 18.3.4.4.4), and the equivalent capability of precise store in prior generation (see Section 18.3.6.3).

In the core PMU of the 4th generation Intel Core processor, load latency facility and PDIR capabilities are unchanged. However, precise store is replaced by an enhanced capability, data address profiling, that is not restricted to store address. Data address profiling also records information in PEBS records at offsets 98H, A0H, and ABH.

### 18.3.6.3 PEBS Data Address Profiling

The Data Linear Address facility is also abbreviated as DataLA. The facility is a replacement or extension of the precise store facility in previous processor generations. The DataLA facility complements the load latency facility by providing a means to profile load and store memory references in the system, leverages the PEBS facility, and provides additional information about sampled loads and stores. Having precise memory reference events with linear address information for both loads and stores provides information to improve data structure layout, eliminate remote node references, and identify cache-line conflicts in NUMA systems.

The DataLA facility in the 4th generation processor supports the following events configured to use PEBS:

**Table 18-25. Precise Events That Supports Data Linear Address Profiling**

| Event Name                       | Event Name                                |
|----------------------------------|---|
| MEM_UOPS_RETIRED.STLB_MISS_LOADS | MEM_UOPS_RETIRED.STLB_MISS_STORES         |
| MEM_UOPS_RETIRED.LOCK_LOADS      | MEM_UOPS_RETIRED.SPLIT_STORES             |
| MEM_UOPS_RETIRED.SPLIT_LOADS     | MEM_UOPS_RETIRED.ALL_STORES               |
| MEM_UOPS_RETIRED.ALL_LOADS       | MEM_LOAD_UOPS_LLC_MISS_RETIRED.LOCAL_DRAM |
| MEM_LOAD_UOPS_RETIRED.L1_HIT     | MEM_LOAD_UOPS_RETIRED.L2_HIT              |
| MEM_LOAD_UOPS_RETIRED.L3_HIT     | MEM_LOAD_UOPS_RETIRED.L1_MISS             |
| MEM_LOAD_UOPS_RETIRED.L2_MISS    | MEM_LOAD_UOPS_RETIRED.L3_MISS             |
| MEM_LOAD_UOPS_RETIRED.HIT_LFB    | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_MISS    |

**Table 18-25. Precise Events That Supports Data Linear Address Profiling (Contd.)**

| Event Name                                    | Event Name                              |
|---|---|
| MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HIT         | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HITM  |
| UOPS_RETIRED.ALL (if load or store is tagged) | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_NONE |

DataLA can use any one of the IA32\_PMC0-IA32\_PMC3 counters. Counter overflows will initiate the generation of PEBS records. Upon counter overflow, hardware captures the linear address and possible other status information of the retiring memory uop. This information is then written to the PEBS record that is subsequently generated.

To enable the DataLA facility, software must complete the following steps. Please note that the DataLA facility relies on the PEBS facility, so the PEBS configuration requirements must be completed before attempting to capture DataLA information.

- Complete the PEBS configuration steps.
- Program the an event listed in Table 18-25 using any one of IA32\_PERFEVTSEL0-IA32\_PERFEVTSEL3.
- Set the corresponding IA32\_PEBS\_ENABLE.PEBS\_EN\_CTRx bit. This enables the corresponding IA32\_PMCx as a PEBS counter and enables the DataLA facility.

When the DataLA facility is enabled, the relevant information written into a PEBS record affects entries at offsets 98H, A0H and A8H, as shown in Table 18-26.

**Table 18-26. Layout of Data Linear Address Information In PEBS Record**

| Field               | Offset | Description   |
|---------------------|--------|---|
| Data Linear Address | 98H    | The linear address of the load or the destination of the store.   |
| Store Status        | A0H    | <ul style="list-style-type: none"> <li>▪ <b>DCU Hit</b> (Bit 0): The store hit the data cache closest to the core (L1 cache) if this bit is set, otherwise the store missed the data cache. This information is valid only for the following store events: UOPS_RETIRED.ALL (if store is tagged), MEM_UOPS_RETIRED.STLB_MISS_STORES, MEM_UOPS_RETIRED.SPLIT_STORES, MEM_UOPS_RETIRED.ALL_STORES</li> <li>▪ Other bits are zero, The STLB_MISS, LOCK bit information can be obtained by programming the corresponding store event in Table 18-25.</li> </ul> |
| Reserved            | A8H    | Always zero.  |

### 18.3.6.3.1 EventingIP Record

The PEBS record layout for processors based on Intel microarchitecture code name Haswell adds a new field at offset 0B0H. This is the eventingIP field that records the IP address of the retired instruction that triggered the PEBS assist. The EIP/RIP field at offset 08H records the IP address of the next instruction to be executed following the PEBS assist.

### 18.3.6.4 Off-core Response Performance Monitoring

The core PMU facility to collect off-core response events are similar to those described in Section 18.3.4.5. The event codes are listed in Table 18-15. Each event code for off-core response monitoring requires programming an associated configuration MSR, MSR\_OFFCORE\_RSP\_x. Software must program MSR\_OFFCORE\_RSP\_x according to:

- Transaction request type encoding (bits 15:0): see Table 18-27.
- Supplier information (bits 30:16): see Table 18-28.
- Snoop response information (bits 37:31): see Table 18-18.

**Table 18-27. MSR\_OFFCORE\_RSP\_x Request\_Type Definition (Haswell microarchitecture)**

| Bit Name           | Offset | Description  |
|--------------------|--------|--|
| DMND_DATA_RD       | 0      | (R/W). Counts the number of demand data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count L2 data read prefetches or instruction fetches. |
| DMND_RFO           | 1      | (R/W). Counts the number of demand and DCU prefetch reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO prefetches.                                     |
| DMND_IFETCH        | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.   |
| COREWB             | 3      | (R/W). Counts the number of modified cachelines written back.  |
| PF_DATA_RD         | 4      | (R/W). Counts the number of data cacheline reads generated by L2 prefetchers.  |
| PF_RFO             | 5      | (R/W). Counts the number of RFO requests generated by L2 prefetchers.  |
| PF_IFETCH          | 6      | (R/W). Counts the number of code reads generated by L2 prefetchers.  |
| PF_L3_DATA_RD      | 7      | (R/W). Counts the number of data cacheline reads generated by L3 prefetchers.  |
| PF_L3_RFO          | 8      | (R/W). Counts the number of RFO requests generated by L3 prefetchers.  |
| PF_L3_CODE_RD      | 9      | (R/W). Counts the number of code reads generated by L3 prefetchers.  |
| SPLIT_LOCK_UC_LOCK | 10     | (R/W). Counts the number of lock requests that split across two cachelines or are to UC memory.  |
| STRM_ST            | 11     | (R/W). Counts the number of streaming store requests electronically.   |
| Reserved           | 12-14  | Reserved   |
| OTHER              | 15     | (R/W). Any other request that crosses IDI, including I/O.  |

The supplier information field listed in Table 18-28. The fields vary across products (according to CPUID signatures) and is noted in the description.

**Table 18-28. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definition (CPUID Signature 06\_3CH, 06\_46H)**

| Subtype       | Bit Name | Offset | Description                                    |
|---------------|----------|--------|--|
| Common        | Any      | 16     | (R/W). Catch all value for any response types. |
| Supplier Info | NO_SUPP  | 17     | (R/W). No Supplier Information available       |
|               | L3_HITM  | 18     | (R/W). M-state initial lookup stat in L3.      |
|               | L3_HITE  | 19     | (R/W). E-state                                 |
|               | L3_HITS  | 20     | (R/W). S-state                                 |
|               | Reserved | 21     | Reserved                                       |
|               | LOCAL    | 22     | (R/W). Local DRAM Controller                   |
|               | Reserved | 30:23  | Reserved                                       |

**Table 18-29. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definition (CPUID Signature 06\_45H)**

| Subtype       | Bit Name               | Offset | Description                                    |
|---------------|------------------------|--------|--|
| Common        | Any                    | 16     | (R/W). Catch all value for any response types. |
| Supplier Info | NO_SUPP                | 17     | (R/W). No Supplier Information available       |
|               | L3_HITM                | 18     | (R/W). M-state initial lookup stat in L3.      |
|               | L3_HITE                | 19     | (R/W). E-state                                 |
|               | L3_HITS                | 20     | (R/W). S-state                                 |
|               | Reserved               | 21     | Reserved                                       |
|               | L4_HIT_LOCAL_L4        | 22     | (R/W). L4 Cache                                |
|               | L4_HIT_REMOTE_HOP0_L4  | 23     | (R/W). L4 Cache                                |
|               | L4_HIT_REMOTE_HOP1_L4  | 24     | (R/W). L4 Cache                                |
|               | L4_HIT_REMOTE_HOP2P_L4 | 25     | (R/W). L4 Cache                                |
|               | Reserved               | 30:26  | Reserved                                       |

#### 18.3.6.4.1 Off-core Response Performance Monitoring in Intel Xeon Processors E5 v3 Series

Table 18-28 lists the supplier information field that apply to Intel Xeon processor E5 v3 series (CPUID signature 06\_3FH).

**Table 18-30. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definition**

| Subtype       | Bit Name             | Offset | Description                                    |
|---------------|----------------------|--------|--|
| Common        | Any                  | 16     | (R/W). Catch all value for any response types. |
| Supplier Info | NO_SUPP              | 17     | (R/W). No Supplier Information available       |
|               | L3_HITM              | 18     | (R/W). M-state initial lookup stat in L3.      |
|               | L3_HITE              | 19     | (R/W). E-state                                 |
|               | L3_HITS              | 20     | (R/W). S-state                                 |
|               | L3_HITF              | 21     | (R/W). F-state                                 |
|               | LOCAL                | 22     | (R/W). Local DRAM Controller                   |
|               | Reserved             | 26:23  | Reserved                                       |
|               | L3_MISS_REMOTE_HOP0  | 27     | (R/W). Hop 0 Remote supplier                   |
|               | L3_MISS_REMOTE_HOP1  | 28     | (R/W). Hop 1 Remote supplier                   |
|               | L3_MISS_REMOTE_HOP2P | 29     | (R/W). Hop 2 or more Remote supplier           |
|               | Reserved             | 30     | Reserved                                       |

#### 18.3.6.5 Performance Monitoring and Intel® TSX


Chapter 16 of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1* describes the details of Intel® Transactional Synchronization Extensions (Intel TSX). This section describes performance monitoring support for Intel TSX.

If a processor supports Intel TSX, the core PMU enhances its IA32\_PERFEVTSELx MSR with two additional bit fields for event filtering. Support for Intel TSX is indicated by either (a) CPUID.(EAX=7, ECX=0):RTM[bit 11]=1, or (b) if CPUID.07H.EBX.HLE [bit 4] = 1. The TSX-enhanced layout of IA32\_PERFEVTSELx is shown in Figure 18-34. The two additional bit fields are:

- **IN\_TX** (bit 32): When set, the counter will only include counts that occurred inside a transactional region, regardless of whether that region was aborted or committed. This bit may only be set if the processor supports HLE or RTM.
- **IN\_TXCP** (bit 33): When set, the counter will not include counts that occurred inside of an aborted transactional region. This bit may only be set if the processor supports HLE or RTM. This bit may only be set for IA32\_PERFEVTSEL2.

When the IA32\_PERFEVTSELx MSR is programmed with both IN\_TX=0 and IN\_TXCP=0 on a processor that supports Intel TSX, the result in a counter may include detectable conditions associated with a transaction code region for its aborted execution (if any) and completed execution.

In the initial implementation, software may need to take pre-caution when using the IN\_TXCP bit. see Table 2-28.



**Figure 18-34. Layout of IA32\_PERFEVTSELx MSRs Supporting Intel TSX**

A common usage of setting IN\_TXCP=1 is to capture the number of events that were discarded due to a transactional abort. With IA32\_PMC2 configured to count in such a manner, then when a transactional region aborts, the value for that counter is restored to the value it had prior to the aborted transactional region. As a result, any updates performed to the counter during the aborted transactional region are discarded.

On the other hand, setting IN\_TX=1 can be used to drill down on the performance characteristics of transactional code regions. When a PMCx is configured with the corresponding IA32\_PERFEVTSELx.IN\_TX=1, only eventing conditions that occur inside transactional code regions are propagated to the event logic and reflected in the counter result. Eventing conditions specified by IA32\_PERFEVTSELx but occurring outside a transactional region are discarded. The following example illustrates using three counters to drill down cycles spent inside and outside of transactional regions:

- Program IA32\_PERFEVTSEL2 to count Unhalted\_Core\_Cycles with (IN\_TXCP=1, IN\_TX=0), such that IA32\_PMC2 will count cycles spent due to aborted TSX transactions;
- Program IA32\_PERFEVTSEL0 to count Unhalted\_Core\_Cycles with (IN\_TXCP=0, IN\_TX=1), such that IA32\_PMC0 will count cycles spent by the transactional code regions;
- Program IA32\_PERFEVTSEL1 to count Unhalted\_Core\_Cycles with (IN\_TXCP=0, IN\_TX=0), such that IA32\_PMC1 will count total cycles spent by the non-transactional code and transactional code regions.

Additionally, a number of performance events are solely focused on characterizing the execution of Intel TSX transactional code, they are listed in Table 19-10.

#### 18.3.6.5.1 Intel TSX and PEBS Support

If a PEBS event would have occurred inside a transactional region, then the transactional region first aborts, and then the PEBS event is processed.

Two of the TSX performance monitoring events in Table 19-10 also support using PEBS facility to capture additional information. They are:

- HLE\_RETIREDA.BORT ED (encoding C8H mask 04H),
- RTM\_RETIREDA.BORTED (encoding C9H mask 04H).

A transactional abort (HLE\_RETIREDA.BORTED,RTM\_RETIREDA.BORTED) can also be programmed to cause PEBS events. In this scenario, a PEBS event is processed following the abort.

Pending a PEBS record inside of a transactional region will cause a transactional abort. If a PEBS record was pended at the time of the abort or on an overflow of the TSX PEBS events listed above, only the following PEBS entries will be valid (enumerated by PEBS entry offset B8H bits[33:32] to indicate an HLE abort or an RTM abort):

- Offset B0H: EventingIP,
- Offset B8H: TX Abort Information

These fields are set for all PEBS events.

- Offset 08H (RIP/EIP) corresponds to the instruction following the outermost XACQUIRE in HLE or the first instruction of the fallback handler of the outermost XBEGIN instruction in RTM. This is useful to identify the aborted transactional region.

In the case of HLE, an aborted transaction will restart execution deterministically at the start of the HLE region. In the case of RTM, an aborted transaction will transfer execution to the RTM fallback handler.

The layout of the TX Abort Information field is given in Table 18-31.

**Table 18-31. TX Abort Information Field Definition**

| Bit Name              | Offset | Description   |
|-----------------------|--------|---|
| Cycles_Last_TX        | 31:0   | The number of cycles in the last TSX region, regardless of whether that region had aborted or committed.                              |
| HLE_Abort             | 32     | If set, the abort information corresponds to an aborted HLE execution   |
| RTM_Abort             | 33     | If set, the abort information corresponds to an aborted RTM execution   |
| Instruction_Abort     | 34     | If set, the abort was associated with the instruction corresponding to the eventing IP (offset 0B0H) within the transactional region. |
| Non_Instruction_Abort | 35     | If set, the instruction corresponding to the eventing IP may not necessarily be related to the transactional abort.                   |
| Retry                 | 36     | If set, retrying the transactional execution may have succeeded.  |
| Data_Conflict         | 37     | If set, another logical processor conflicted with a memory address that was part of the transactional region that aborted.            |
| Capacity Writes       | 38     | If set, the transactional region aborted due to exceeding resources for transactional writes.   |
| Capacity Reads        | 39     | If set, the transactional region aborted due to exceeding resources for transactional reads.  |
| Reserved              | 63:40  | Reserved  |

### 18.3.6.6 Uncore Performance Monitoring Facilities in the 4th Generation Intel® Core™ Processors

The uncore sub-system in the 4th Generation Intel® Core™ processors provides its own performance monitoring facility. The uncore PMU facility provides dedicated MSRs to select uncore performance monitoring events in a similar manner as those described in Section 18.3.4.6.

The ARB unit and each C-Box provide local pairs of event select MSR and counter register. The layout of the event select MSRs in the C-Boxes are identical as shown in Figure 18-32.

At the uncore domain level, there is a master set of control MSRs that centrally manages all the performance monitoring facility of uncore units. Figure 18-33 shows the layout of the uncore domain global control.

Additionally, there is also a fixed counter, counting uncore clockticks, for the uncore domain. Table 18-19 summarizes the number MSRs for uncore PMU for each box.

**Table 18-32. Uncore PMU MSR Summary**

| Box           | # of Boxes   | Counters per Box | Counter Width | General Purpose | Global Enable | Comment                                       |
|---------------|--------------|------------------|---------------|-----------------|---------------|---|
| C-Box         | SKU specific | 2                | 44            | Yes             | Per-box       | Up to 4, see Table 2-20<br>MSR_UNC_CBO_CONFIG |
| ARB           | 1            | 2                | 44            | Yes             | Uncore        |   |
| Fixed Counter | N.A.         | N.A.             | 48            | No              | Uncore        |   |

The uncore performance events for the C-Box and ARB units are listed in Table 19-11.

### 18.3.6.7 Intel® Xeon® Processor E5 v3 Family Uncore Performance Monitoring Facility

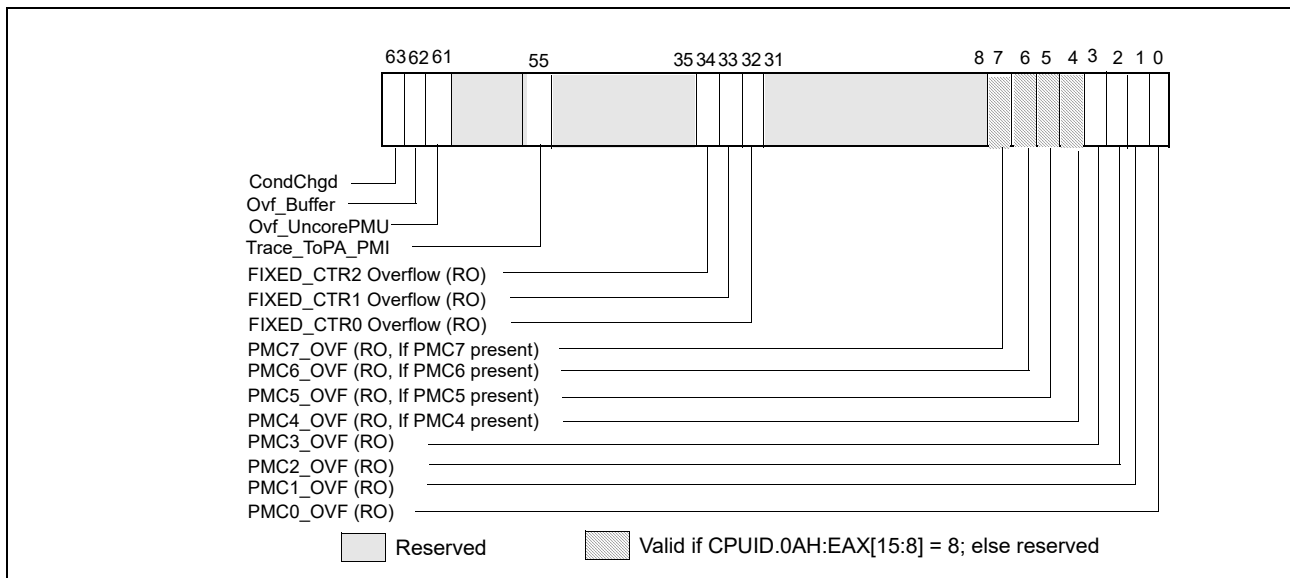
Details of the uncore performance monitoring facility of Intel Xeon Processor E5 v3 families are available in “Intel® Xeon® Processor E5 v3 Uncore Performance Monitoring Programming Reference Manual”. The MSR-based uncore PMU interfaces are listed in Table 2-32.

### 18.3.7 5th Generation Intel® Core™ Processor and Intel® Core™ M Processor Performance Monitoring Facility

The 5th Generation Intel® Core™ processor and the Intel® Core™ M processor families are based on the Broadwell microarchitecture. The core PMU supports architectural performance monitoring capability with version ID 3 (see Section 18.2.3) and a host of non-architectural monitoring capabilities.

Architectural performance monitoring version 3 capabilities are described in Section 18.2.3.

The core PMU has the same capability as those described in Section 18.3.6. IA32\_PERF\_GLOBAL\_STATUS provide a bit indicator (bit 55) for PMI handler to distinguish PMI due to output buffer overflow condition due to accumulating packet data from Intel Processor Trace.



**Figure 18-35. IA32\_PERF\_GLOBAL\_STATUS MSR in Broadwell Microarchitecture**

Details of Intel Processor Trace is described in Chapter 35, “Intel® Processor Trace”. IA32\_PERF\_GLOBAL\_OVF\_CTRL MSR provide a corresponding reset control bit.




Figure 18-36. IA32\_PERF\_GLOBAL\_OVF\_CTRL MSR in Broadwell microarchitecture

The specifics of non-architectural performance events are listed in Chapter 19, “Performance Monitoring Events”.

### 18.3.8 6th Generation, 7th Generation and 8th Generation Intel® Core™ Processor Performance Monitoring Facility

The 6th generation Intel® Core™ processor is based on the Skylake microarchitecture. The 7th generation Intel® Core™ processor is based on the Kaby Lake microarchitecture. The 8th generation Intel® Core™ processor is based on the Coffee Lake microarchitecture. For these microarchitectures, the core PMU supports architectural performance monitoring capability with version ID 4 (see Section 18.2.4) and a host of non-architectural monitoring capabilities.

Architectural performance monitoring version 4 capabilities are described in Section 18.2.4.

The core PMU’s capability is similar to those described in Section 18.6.3 through Section 18.3.4.5, with some differences and enhancements summarized in Table 18-22. Additionally, the core PMU provides some enhancement to support performance monitoring when the target workload contains instruction streams using Intel® Transactional Synchronization Extensions (TSX), see Section 18.3.6.5. For details of Intel TSX, see Chapter 16, “Programming with Intel® Transactional Synchronization Extensions” of *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1*.

Performance monitoring result may be affected by side-band activity on processors that support Intel SGX, details are described in Chapter 42, “Enclave Code Debug and Profiling”.



Table 18-33. Core PMU Comparison

| Box  | Intel® Microarchitecture Code Name<br>Skylake, Kaby Lake and Coffee Lake  | Intel® Microarchitecture Code<br>Name Haswell and Broadwell  | Comment  |
|--|---|--|--|
| # of Fixed counters per thread                                       | 3   | 3  |  |
| # of general-purpose counters per core                               | 8   | 8  |  |
| Counter width (R,W)  | R:48, W: 32/48  | R:48, W: 32/48   | See Section 18.2.2.  |
| # of programmable counters per thread                                | 4 or (8 if a core not shared by two threads)  | 4 or (8 if a core not shared by two threads)   | CPUID enumerates # of counters.  |
| Architectural Perfmon version  | 4   | 3  | See Section 18.2.4   |
| PMI Overhead Mitigation  | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with streamlined semantics.</li> <li>▪ Freeze_on_LBR with streamlined semantics.</li> <li>▪ Freeze_while_SMM.</li> </ul> | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_on_LBR with legacy semantics for branch profiling.</li> <li>▪ Freeze_while_SMM.</li> </ul> | See Section 17.4.7. Legacy semantics not supported with version 4 or higher. |
| Counter and Buffer Overflow Status Management                        | <ul style="list-style-type: none"> <li>▪ Query via IA32_PERF_GLOBAL_STATUS</li> <li>▪ Reset via IA32_PERF_GLOBAL_STATUS_RESET</li> <li>▪ Set via IA32_PERF_GLOBAL_STATUS_SET</li> </ul> | <ul style="list-style-type: none"> <li>▪ Query via IA32_PERF_GLOBAL_STATUS</li> <li>▪ Reset via IA32_PERF_GLOBAL_OVF_CTRL</li> </ul>   | See Section 18.2.4.  |
| IA32_PERF_GLOBAL_STATUS Indicators of Overflow/Overhead/Interference | <ul style="list-style-type: none"> <li>▪ Individual counter overflow</li> <li>▪ PEBS buffer overflow</li> <li>▪ ToPA buffer overflow</li> <li>▪ CTR_Frz, LBR_Frz, ASCI</li> </ul>       | <ul style="list-style-type: none"> <li>▪ Individual counter overflow</li> <li>▪ PEBS buffer overflow</li> <li>▪ ToPA buffer overflow (applicable to Broadwell microarchitecture)</li> </ul>        | See Section 18.2.4.  |
| Enable control in IA32_PERF_GLOBAL_STATUS                            | <ul style="list-style-type: none"> <li>▪ CTR_Frz</li> <li>▪ LBR_Frz</li> </ul>  | NA   | See Section 18.2.4.1.  |
| Perfmon Counter In-Use Indicator                                     | Query IA32_PERF_GLOBAL_INUSE  | NA   | See Section 18.2.4.3.  |
| Precise Events   | See Table 18-36.  | See Table 18-12.   | IA32_PMC4-PMC7 do not support PEBS.  |
| PEBS for front end events  | See Section 18.3.8.1.4.   | No   |  |
| LBR Record Format Encoding   | 000101b   | 000100b  | Section 17.4.8.1   |
| LBR Size   | 32 entries  | 16 entries   |  |
| LBR Entry  | From_IP/To_IP/LBR_Info triplet  | From_IP/To_IP pair   | Section 17.12  |
| LBR Timing   | Yes   | No   | Section 17.12.1  |
| Call Stack Profiling   | Yes, see Section 17.11  | Yes, see Section 17.11   | Use LBR facility   |
| Off-core Response Event  | MSR 1A6H and 1A7H; Extended request and response types.   | MSR 1A6H and 1A7H; Extended request and response types.  |  |
| Intel TSX support for Perfmon  | See Section 18.3.6.5.   | See Section 18.3.6.5.  |  |

### 18.3.8.1 Processor Event Based Sampling (PEBS) Facility

The PEBS facility in the 6th generation, 7th generation and 8th generation Intel Core processors provides a number enhancement relative to PEBS in processors based on Haswell/Broadwell microarchitectures. The key components and differences of PEBS facility relative to Haswell/Broadwell microarchitecture is summarized in Table 18-34.

**Table 18-34. PEBS Facility Comparison**

| Box                           | Intel® Microarchitecture Code Name Skylake, Kaby Lake and Coffee Lake          | Intel® Microarchitecture Code Name Haswell and Broadwell    | Comment                                  |
|-------------------------------|--|---|--|
| Valid IA32_PMCx               | PMCO-PMC3  | PMCO-PMC3   | No PEBS on PMC4-PMC7.                    |
| PEBS Buffer Programming       | Section 18.3.1.1.1   | Section 18.3.1.1.1  | Unchanged                                |
| IA32_PEBS_ENABLE Layout       | Figure 18-15   | Figure 18-15  |  |
| PEBS-EventingIP               | Yes  | Yes   |  |
| PEBS record format encoding   | 0011b  | 0010b   |  |
| PEBS record layout            | Table 18-35; enhanced fields at offsets 98H- B8H; and TSC record field at COH. | Table 18-24; enhanced fields at offsets 98H, A0H, A8H, B0H. |  |
| Multi-counter PEBS resolution | PEBS record 90H resolves the eventing counter overflow.                        | PEBS record 90H reflects IA32_PERF_GLOBAL_STATUS.           |  |
| Precise Events                | See Table 18-36.   | See Table 18-12.  | IA32_PMC4-IA32_PMC7 do not support PEBS. |
| PEBS-PDIR                     | Yes  | Yes   | IA32_PMC1 only.                          |
| PEBS-Load Latency             | See Section 18.3.4.4.2.  | See Section 18.3.4.4.2.                                     |  |
| Data Address Profiling        | Yes  | Yes   |  |
| FrontEnd event support        | FrontEnd_Retried event and MSR_PEBS_FRONTEND.                                  | No  | IA32_PMC0-PMC3 only.                     |

Only IA32\_PMC0 through IA32\_PMC3 support PEBS.

**NOTES**

Precise events are only valid when the following fields of IA32\_PERFEVTSELx are all zero: AnyThread, Edge, Invert, CMask.

In a PMU with PDIR capability, PEBS behavior is unpredictable if IA32\_PERFEVTSELx or IA32\_PMCx is changed for a PEBS-enabled counter while an event is being counted. To avoid this, changes to the programming or value of a PEBS-enabled counter should be performed when the counter is disabled.

**18.3.8.1.1 PEBS Data Format**

The PEBS record format for the 6th generation, 7th generation and 8th generation Intel Core processors is reporting with encoding 0011b in IA32\_PERF\_CAPABILITIES[11:8]. The lay out is shown in Table 18-35. The PEBS record format, along with debug/store area storage format, does not change regardless of whether IA-32e mode is active or not. CPUID.01H:ECX.DTES64[bit 2] reports whether the processor's DS storage format support is mode-independent. When set, it uses 64-bit DS storage format.

**Table 18-35. PEBS Record Format for 6th Generation, 7th Generation and 8th Generation Intel Core Processor Families**

| Byte Offset | Field    | Byte Offset | Field                                     |
|-------------|----------|-------------|---|
| 00H         | R/EFLAGS | 68H         | R11                                       |
| 08H         | R/EIP    | 70H         | R12                                       |
| 10H         | R/EAX    | 78H         | R13                                       |
| 18H         | R/EBX    | 80H         | R14                                       |
| 20H         | R/ECX    | 88H         | R15                                       |
| 28H         | R/EDX    | 90H         | Applicable Counter                        |
| 30H         | R/ESI    | 98H         | Data Linear Address                       |
| 38H         | R/EDI    | A0H         | Data Source Encoding                      |
| 40H         | R/EBP    | A8H         | Latency value (core cycles)               |
| 48H         | R/ESP    | B0H         | EventingIP                                |
| 50H         | R8       | B8H         | TX Abort Information (Section 18.3.6.5.1) |
| 58H         | R9       | C0H         | TSC                                       |
| 60H         | R10      |             |   |

The layout of PEBS records are largely identical to those shown in Table 18-24.

The PEBS records at offsets 98H, A0H, and ABH record data gathered from three of the PEBS capabilities in prior processor generations: load latency facility (Section 18.3.4.4.2), PDIR (Section 18.3.4.4.4), and data address profiling (Section 18.3.6.3).

In the core PMU of the 6th generation, 7th generation and 8th generation Intel Core processors, load latency facility and PDIR capabilities and data address profiling are unchanged relative to the 4th generation and 5th generation Intel Core processors. Similarly, precise store is replaced by data address profiling.

With format 0010b, a snapshot of the IA32\_PERF\_GLOBAL\_STATUS may be useful to resolve the situations when more than one of IA32\_PMICx have been configured to collect PEBS data and two consecutive overflows of the PEBS-enabled counters are sufficiently far apart in time. It is also possible for the image at 90H to indicate multiple PEBS-enabled counters have overflowed. In the latter scenario, software cannot to correlate the PEBS record entry to the multiple overflowed bits.

With PEBS record format encoding 0011b, offset 90H reports the “applicable counter” field, which is a multi-counter PEBS resolution index allowing software to correlate the PEBS record entry with the eventing PEBS overflow when multiple counters are configured to record PEBS records. Additionally, offset C0H captures a snapshot of the TSC that provides a time line annotation for each PEBS record entry.

### 18.3.8.1.2 PEBS Events

The list of precise events supported for PEBS in the Skylake, Kaby Lake and Coffee Lake microarchitectures is shown in Table 18-36.

**Table 18-36. Precise Events for the Skylake, Kaby Lake and Coffee Lake Microarchitectures**

| Event Name                           | Event Select | Sub-event                    | UMask |
|--------------------------------------|--------------|------------------------------|-------|
| INST_RETIRED                         | C0H          | PREC_DIST <sup>1</sup>       | 01H   |
|                                      |              | ALL_CYCLES <sup>2</sup>      | 01H   |
| OTHER_ASSISTS                        | C1H          | ANY                          | 3FH   |
| BR_INST_RETIRED                      | C4H          | CONDITIONAL                  | 01H   |
|                                      |              | NEAR_CALL                    | 02H   |
|                                      |              | ALL_BRANCHES                 | 04H   |
|                                      |              | NEAR_RETURN                  | 08H   |
|                                      |              | NEAR_TAKEN                   | 20H   |
|                                      |              | FAR_BRACHES                  | 40H   |
| BR_MISP_RETIRED                      | C5H          | CONDITIONAL                  | 01H   |
|                                      |              | ALL_BRANCHES                 | 04H   |
|                                      |              | NEAR_TAKEN                   | 20H   |
| FRONTEND_RETIRED                     | C6H          | <Programmable <sup>3</sup> > | 01H   |
| HLE_RETIRED                          | C8H          | ABORTED                      | 04H   |
| RTM_RETIRED                          | C9H          | ABORTED                      | 04H   |
| MEM_INST_RETIRED <sup>2</sup>        | D0H          | LOCK_LOADS                   | 21H   |
|                                      |              | SPLIT_LOADS                  | 41H   |
|                                      |              | SPLIT_STORES                 | 42H   |
|                                      |              | ALL_LOADS                    | 81H   |
|                                      |              | ALL_STORES                   | 82H   |
| MEM_LOAD_RETIRED <sup>4</sup>        | D1H          | L1_HIT                       | 01H   |
|                                      |              | L2_HIT                       | 02H   |
|                                      |              | L3_HIT                       | 04H   |
|                                      |              | L1_MISS                      | 08H   |
|                                      |              | L2_MISS                      | 10H   |
|                                      |              | L3_MISS                      | 20H   |
|                                      |              | HIT_LFB                      | 40H   |
| MEM_LOAD_L3_HIT_RETIRED <sup>2</sup> | D2H          | XSNP_MISS                    | 01H   |
|                                      |              | XSNP_HIT                     | 02H   |
|                                      |              | XSNP_HITM                    | 04H   |
|                                      |              | XSNP_NONE                    | 08H   |

**NOTES:**

1. Only available on IA32\_PMC1.
2. INST\_RETIRED.ALL\_CYCLES is configured with additional parameters of cmask = 10 and INV = 1
3. Subevents are specified using MSR\_PEBS\_FRONTEND, see Section 18.3.8.2
4. Instruction with at least one load up experiencing the condition specified in the UMask.

**18.3.8.1.3 Data Address Profiling**

The PEBS Data address profiling on the 6th generation, 7th generation and 8th generation Intel Core processors is largely unchanged from the prior generation. When the DataLA facility is enabled, the relevant information written into a PEBS record affects entries at offsets 98H, A0H and A8H, as shown in Table 18-26.

**Table 18-37. Layout of Data Linear Address Information In PEBS Record**

| Field               | Offset | Description   |
|---------------------|--------|---|
| Data Linear Address | 98H    | The linear address of the load or the destination of the store.   |
| Store Status        | A0H    | <ul style="list-style-type: none"> <li>▪ <b>DCU Hit</b> (Bit 0): The store hit the data cache closest to the core (L1 cache) if this bit is set, otherwise the store missed the data cache. This information is valid only for the following store events: UOPS_RETIRED.ALL (if store is tagged), MEM_INST_RETIRED.STLB_MISS_STORES, MEM_INST_RETIRED.ALL_STORES, MEM_INST_RETIRED.SPLIT_STORES.</li> <li>▪ Other bits are zero.</li> </ul> |
| Reserved            | A8H    | Always zero.  |

#### 18.3.8.1.4 PEBS Facility for Front End Events

In the 6th generation, 7th generation and 8th generation Intel Core processors, the PEBS facility has been extended to allow capturing PEBS data for some microarchitectural conditions related to front end events. The frontend microarchitectural conditions supported by PEBS requires the following interfaces:

- The IA32\_PERFEVTSELx MSR must select "FrontEnd\_Retired" (C6H) in the EventSelect field (bits 7:0) and umask = 01H,
- The "FRONTEND\_RETIRED" event employs a new MSR, MSR\_PEBS\_FRONTEND, to specify the supported frontend event details, see Table 18-38.
- Program the PEBS\_EN\_PMCx field of IA32\_PEBS\_ENABLE MSR as required.

Note the AnyThread field of IA32\_PERFEVTSELx is ignored by the processor for the "FRONTEND\_RETIRED" event.

The sub-event encodings supported by MSR\_PEBS\_FRONTEND.EVTSEL is given in Table 18-38.

**Table 18-38. FrontEnd\_Retired Sub-Event Encodings Supported by MSR\_PEBS\_FRONTEND.EVTSEL**

| Sub-Event Name   | EVTSEL | Description   |
|------------------|--------|---|
| DSB_MISS         | 11H    | Retired Instructions which experienced decode stream buffer (DSB) miss.   |
| L11_MISS         | 12H    | The fetch of retired Instructions which experienced Instruction L1 Cache true miss <sup>1</sup> . Additional requests to the same cache line as an in-flight L11 cache miss will not be counted.  |
| L2_MISS          | 13H    | The fetch of retired Instructions which experienced L2 Cache true miss. Additional requests to the same cache line as an in-flight MLC cache miss will not be counted.  |
| ITLB_MISS        | 14H    | The fetch of retired Instructions which experienced ITLB true miss. Additional requests to the same cache line as an in-flight ITLB miss will not be counted.   |
| STLB_MISS        | 15H    | The fetch of retired Instructions which experienced STLB true miss. Additional requests to the same cache line as an in-flight STLB miss will not be counted.   |
| IDQ_READ_BUBBLES | 6H     | <p>An IDQ read bubble is defined as any one of the 4 allocation slots of IDQ that is not filled by the front-end on any cycle where there is no back end stall. Using the threshold and latency fields in MSR_PEBS_FRONTEND allows counting of IDQ read bubbles of various magnitude and duration. Latency controls the number of cycles and Threshold controls the number of allocation slots that contain bubbles.</p> <p>The event counts if and only if a sequence of at least FE_LATENCY consecutive cycles contain at least FE_TRESHOLD number of bubbles each.</p> |

#### NOTES:

1. A true miss is the first miss for a cacheline/page (excluding secondary misses that fall into same cacheline/page).

The layout of MSR\_PEBS\_FRONTEND is given in Table 18-39.

**Table 18-39. MSR\_PEBS\_FRONTEND Layout**

| Bit Name          | Offset | Description   |
|-------------------|--------|---|
| EVTSEL            | 7:0    | Encodes the sub-event within FrontEnd_Retired that can use PEBS facility, see Table 18-38.                                      |
| IDQ_Bubble_Length | 19:8   | Specifies the threshold of continuously elapsed cycles for the specified width of bubbles when counting IDQ_READ_BUBBLES event. |
| IDQ_Bubble_Width  | 22:20  | Specifies the threshold of simultaneous bubbles when counting IDQ_READ_BUBBLES event.   |
| Reserved          | 63:23  | Reserved  |

**18.3.8.1.5 FRONTEND\_RETIRED**

The FRONTEND\_RETIRED event is designed to help software developers identify exact instructions that caused front-end issues. There are some instances in which the event will, by design, the under-counting scenarios include the following:

- The event counts only retired (non-speculative) front-end events, i.e. events from just true program execution path are counted.
- The event will count once per cacheline (at most). If a cacheline contains multiple instructions which caused front-end misses, the count will be only 1 for that line.
- If the multibyte sequence of an instruction spans across two cachelines and causes a miss it will be recorded once. If there were additional misses in the second cacheline, they will not be counted separately.
- If a multi-uop instruction exceeds the allocation width of one cycle, the bubbles associated with these uops will be counted once per that instruction.
- If 2 instructions are fused (macro-fusion), and either of them or both cause front-end misses, it will be counted once for the fused instruction.
- If a front-end (miss) event occurs outside instruction boundary (e.g. due to processor handling of architectural event), it may be reported for the next instruction to retire.

**18.3.8.2 Off-core Response Performance Monitoring**

The core PMU facility to collect off-core response events are similar to those described in Section 18.3.4.5. Each event code for off-core response monitoring requires programming an associated configuration MSR, MSR\_OFFCORE\_RSP\_x. Software must program MSR\_OFFCORE\_RSP\_x according to:

- Transaction request type encoding (bits 15:0): see Table 18-40.
- Supplier information (bits 30:16): see Table 18-41.
- Snoop response information (bits 37:31): see Table 18-42.

**Table 18-40. MSR\_OFFCORE\_RSP\_x Request\_Type Definition (Skylake, Kaby Lake and Coffee Lake Microarchitectures)**

| Bit Name      | Offset | Description   |
|---------------|--------|---|
| DMND_DATA_RD  | 0      | (R/W). Counts the number of demand data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count hw or sw prefetches. |
| DMND_RFO      | 1      | (R/W). Counts the number of demand reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO prefetches.                           |
| DMND_IFETCH   | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.  |
| Reserved      | 6:3    | Reserved  |
| PF_L3_DATA_RD | 7      | (R/W). Counts the number of MLC prefetches into L3.   |
| PF_L3_RFO     | 8      | (R/W). Counts the number of RFO requests generated by MLC prefetches to L3.   |
| Reserved      | 10:9   | Reserved  |
| STRM_ST       | 11     | (R/W). Counts the number of streaming store requests.   |
| Reserved      | 14:12  | Reserved  |
| OTHER         | 15     | (R/W). Any other request that crosses IDI, including I/O.   |

Table 18-41 lists the supplier information field that applies to 6th generation, 7th generation and 8th generation Intel Core processors. (6th generation Intel Core processor CPUID signatures: 06\_4EH, 06\_5EH; 7th generation and 8th generation Intel Core processor CPUID signatures: 06\_8EH, 06\_9EH).

**Table 18-41. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definition (CPUID Signatures 06\_4EH, 06\_5EH and 06\_8EH, 06\_9EH)**

| Subtype       | Bit Name | Offset | Description  |
|---------------|----------|--------|--|
| Common        | Any      | 16     | (R/W). Catch all value for any response types.                     |
| Supplier Info | NO_SUPP  | 17     | (R/W). No Supplier Information available.                          |
|               | L3_HITM  | 18     | (R/W). M-state initial lookup stat in L3.                          |
|               | L3_HITE  | 19     | (R/W). E-state   |
|               | L3_HITS  | 20     | (R/W). S-state   |
|               | Reserved | 21     | Reserved   |
|               | L4_HIT   | 22     | (R/W). L4 Cache (if L4 is present in the processor)                |
|               | Reserved | 25:23  | Reserved   |
|               | DRAM     | 26     | (R/W). Local Node  |
|               | Reserved | 29:27  | Reserved   |
|               | SPL_HIT  | 30     | (R/W). L4 cache super line hit (if L4 is present in the processor) |

Table 18-42 lists the snoop information field that apply to processors with CPUID signatures 06\_4EH, 06\_5EH, 06\_8EH, 06\_9E, and 06\_55H.

**Table 18-42. MSR\_OFFCORE\_RSP\_x Snoop Info Field Definition  
(CPUID Signatures 06\_4EH, 06\_5EH, 06\_8EH, 06\_9E and 06\_55H)**

| Subtype    | Bit Name           | Offset | Description  |
|------------|--------------------|--------|--|
| Snoop Info | SNOOP_NONE         | 31     | (R/W). No details on snoop-related information   |
|            | SNOOP_NOT_NEEDED   | 32     | (R/W). No snoop was needed to satisfy the request.   |
|            | SNOOP_MISS         | 33     | (R/W). A snoop was needed and it missed all snooped caches:<br>-For LLC Hit, ReslHitl was returned by all cores<br>-For LLC Miss, Rspl was returned by all sockets and data was returned from DRAM.  |
|            | SNOOP_HIT_NO_FWD   | 34     | (R/W). A snoop was needed and it hits in at least one snooped cache. Hit denotes a cache-line was valid before snoop effect. This includes:<br>-Snoop Hit w/ Invalidation (LLC Hit, RFO)<br>-Snoop Hit, Left Shared (LLC Hit/Miss, IFetch/Data_RD)<br>-Snoop Hit w/ Invalidation and No Forward (LLC Miss, RFO Hit S)<br>In the LLC Miss case, data is returned from DRAM. |
|            | SNOOP_HIT_WITH_FWD | 35     | (R/W). A snoop was needed and data was forwarded from a remote socket. This includes:<br>-Snoop Forward Clean, Left Shared (LLC Hit/Miss, IFetch/Data_RD/RFT).   |
|            | SNOOP_HITM         | 36     | (R/W). A snoop was needed and it HitM-ed in local or remote cache. HitM denotes a cache-line was in modified state before effect as a results of snoop. This includes:<br>-Snoop HitM w/ WB (LLC miss, IFetch/Data_RD)<br>-Snoop Forward Modified w/ Invalidation (LLC Hit/Miss, RFO)<br>-Snoop MtoS (LLC Hit, IFetch/Data_RD).  |
|            | SNOOP_NON_DRAM     | 37     | (R/W). Target was non-DRAM system address. This includes MMIO transactions.  |

### 18.3.8.2.1 Off-core Response Performance Monitoring for the Intel® Xeon® Processor Scalable Family

The following tables list the requestor and supplier information fields that apply to the Intel® Xeon® Processor Scalable Family.

- Transaction request type encoding (bits 15:0): see Table 18-43.
- Supplier information (bits 30:16): see Table 18-44.
- Snoop response information has not been changed and is the same as in (bits 37:31): see Table 18-42.



**Table 18-43. MSR\_OFFCORE\_RSP\_x Request\_Type Definition (Intel® Xeon® Processor Scalable Family)**

| Bit Name         | Offset | Description   |
|------------------|--------|---|
| DEMAND_DATA_RD   | 0      | (R/W). Counts the number of demand data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count hw or sw prefetches. |
| DEMAND_RFO       | 1      | (R/W). Counts the number of demand reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO prefetches.                           |
| DEMAND_CODE_RD   | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.  |
| Reserved         | 3      | Reserved.   |
| PF_L2_DATA_RD    | 4      | (R/W). Counts the number of prefetch data reads into L2.  |
| PF_L2_RFO        | 5      | (R/W). Counts the number of RFO Requests generated by the MLC prefetches to L2.   |
| Reserved         | 6      | Reserved.   |
| PF_L3_DATA_RD    | 7      | (R/W). Counts the number of MLC data read prefetches into L3.   |
| PF_L3_RFO        | 8      | (R/W). Counts the number of RFO requests generated by MLC prefetches to L3.   |
| Reserved         | 9      | Reserved.   |
| PF_L1D_AND_SW    | 10     | (R/W). Counts data cacheline reads generated by hardware L1 data cache prefetcher or software prefetch requests.  |
| STREAMING_STORES | 11     | (R/W). Counts the number of streaming store requests.   |
| Reserved         | 14:12  | Reserved.   |
| OTHER            | 15     | (R/W). Any other request that crosses IDI, including I/O.   |

Table 18-44 lists the supplier information field that applies to the Intel Xeon Processor Scalable Family (CPUID signature: 06\_55H).

**Table 18-44. MSR\_OFFCORE\_RSP\_x Supplier Info Field Definition (CPUID Signature 06\_55H)**

| Subtype       | Bit Name                  | Offset    | Description   |
|---------------|---------------------------|-----------|---|
| Common        | Any                       | 16        | (R/W). Catch all value for any response types.  |
| Supplier Info | SUPPLIER_NONE             | 17        | (R/W). No Supplier Information available.   |
|               | L3_HIT_M                  | 18        | (R/W). M-state initial lookup stat in L3.   |
|               | L3_HIT_E                  | 19        | (R/W). E-state  |
|               | L3_HIT_S                  | 20        | (R/W). S-state  |
|               | L3_HIT_F                  | 21        | (R/W). F-state  |
|               | Reserved                  | 25:22     | Reserved.   |
|               | L3_MISS_LOCAL_DRAM        | 26        | (R/W). L3 Miss: local home requests that missed the L3 cache and were serviced by local DRAM. |
|               | L3_MISS_REMOTE_HOPO_DRAM  | 27        | (R/W). Hop 0 Remote supplier.   |
|               | L3_MISS_REMOTE_HOP1_DRAM  | 28        | (R/W). Hop 1 Remote supplier.   |
|               | L3_MISS_REMOTE_HOP2P_DRAM | 29        | (R/W). Hop 2 or more Remote supplier.   |
| Reserved      | 30                        | Reserved. |   |

## 18.4 PERFORMANCE MONITORING (INTEL® XEON™ PHI PROCESSORS)

### NOTE

This section also applies to the Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series based on Knights Mill microarchitecture.

### 18.4.1 Intel® Xeon Phi™ Processor 7200/5200/3200 Performance Monitoring

The Intel® Xeon Phi™ processor 7200/5200/3200 series are based on the Knights Landing microarchitecture. The performance monitoring capabilities are distributed between its tiles (pair of processor cores) and untile (connecting many tiles in a physical processor package). Functional details of the tiles and untile of the Knights Landing microarchitecture can be found in Chapter 16 of *Intel® 64 and IA-32 Architectures Optimization Reference Manual*.

A complete description of the tile and untile PMU programming interfaces for Intel Xeon Phi processors based on the Knights Landing microarchitecture can be found in the Technical Document section at <http://www.intel.com/content/www/us/en/processors/xeon/xeon-phi-detail.html>.

A tile contains a pair of cores attached to a shared L2 cache and is similar to those found in Intel® Atom™ processors based on the Silvermont microarchitecture. The processor provides several new capabilities on top of the Silvermont performance monitoring facilities.

The processor supports architectural performance monitoring capability with version ID 3 (see Section 18.2.3) and a host of non-architectural performance monitoring capabilities. The processor provides two general-purpose performance counters (IA32\_PMC0, IA32\_PMC1) and three fixed-function performance counters (IA32\_FIXED\_CTR0, IA32\_FIXED\_CTR1, IA32\_FIXED\_CTR2).

Non-architectural performance monitoring in the processor also uses the IA32\_PERFEVTSELx MSR to configure a set of non-architecture performance monitoring events to be counted by the corresponding general-purpose performance counter.

The bit fields within each IA32\_PERFEVTSELx MSR are defined in Figure 18-6 and described in Section 18.2.1.1 and Section 18.2.3 in the SDM. The processor supports AnyThread counting in three architectural performance monitoring events.

#### 18.4.1.1 Enhancements of Performance Monitoring in the Intel® Xeon Phi™ processor Tile

The Intel® Xeon Phi™ processor tile includes the following enhancements to the Silvermont microarchitecture.

- AnyThread support. This facility is limited to following three architectural events: Instructions Retired, Unhalted Core Cycles, Unhalted Reference Cycles using IA32\_FIXED\_CTR0-2 and Unhalted Core Cycles, Unhalted Reference Cycles using IA32\_PERFEVTSELx.
- PEBS-DLA (Processor Event-Based Sampling-Data Linear Address) fields. The processor provides memory address in addition to the Silvermont PEBS record support on select events. The PEBS recording format as reported by IA32\_PERF\_CAPABILITIES [11:8] is 2.
- Off-core response counting facility. This facility in the processor core allows software to count certain transaction responses between the processor tile to subsystems outside the tile (untile). Counting off-core response requires additional event qualification configuration facility in conjunction with IA32\_PERFEVTSELx. Two off-core response MSRs are provided to use in conjunction with specific event codes that must be specified with IA32\_PERFEVTSELx. Two cores do not share the off-core response MSRs. Knights Landing expands off-core response capability to match the processor untile changes.
- Average request latency measurement. The off-core response counting facility can be combined to use two performance counters to count the occurrences and weighted cycles of transaction requests. This facility is updated to match the processor untile changes.

### 18.4.1.1.1 Processor Event-Based Sampling

The processor supports processor event based sampling (PEBS). PEBS is supported using IA32\_PMC0 (see also Section 17.4.9, “BTS and DS Save Area”).

PEBS uses a debug store mechanism to store a set of architectural state information for the processor. The information provides architectural state of the instruction executed after the instruction that caused the event (See Section 18.6.2.4).

The list of PEBS events supported in the processor is shown in the following table.

**Table 18-45. PEBS Performance Events for the Knights Landing Microarchitecture**

| Event Name       | Event Select | Sub-event           | UMask | Data Linear Address Support |
|------------------|--------------|---------------------|-------|-----------------------------|
| BR_INST_RETIRED  | C4H          | ALL_BRANCHES        | 00H   | No                          |
|                  |              | JCC                 | 7EH   | No                          |
|                  |              | TAKEN_JCC           | FEH   | No                          |
|                  |              | CALL                | F9H   | No                          |
|                  |              | REL_CALL            | FDH   | No                          |
|                  |              | IND_CALL            | FBH   | No                          |
|                  |              | NON_RETURN_IND      | EBH   | No                          |
|                  |              | FAR_BRANCH          | BFH   | No                          |
|                  |              | RETURN              | F7H   | No                          |
| BR_MISP_RETIRED  | C5H          | ALL_BRANCHES        | 00H   | No                          |
|                  |              | JCC                 | 7EH   | No                          |
|                  |              | TAKEN_JCC           | FEH   | No                          |
|                  |              | IND_CALL            | FBH   | No                          |
|                  |              | NON_RETURN_IND      | EBH   | No                          |
|                  |              | RETURN              | F7H   | No                          |
| MEM_UOPS_RETIRED | 04H          | L2_HIT_LOADS        | 02H   | Yes                         |
|                  |              | L2_MISS_LOADS       | 04H   | Yes                         |
|                  |              | DLTB_MISS_LOADS     | 08H   | Yes                         |
| RECYCLEQ         | 03H          | LD_BLOCK_ST_FORWARD | 01H   | Yes                         |
|                  |              | LD_SPLITS           | 08H   | Yes                         |

The PEBS record format 2 supported by processors based on the Knights Landing microarchitecture is shown in Table 18-46, and each field in the PEBS record is 64 bits long.

**Table 18-46. PEBS Record Format for the Knights Landing Microarchitecture**

| Byte Offset | Field    | Byte Offset | Field                   |
|-------------|----------|-------------|-------------------------|
| 00H         | R/EFLAGS | 60H         | R10                     |
| 08H         | R/EIP    | 68H         | R11                     |
| 10H         | R/EAX    | 70H         | R12                     |
| 18H         | R/EBX    | 78H         | R13                     |
| 20H         | R/ECX    | 80H         | R14                     |
| 28H         | R/EDX    | 88H         | R15                     |
| 30H         | R/ESI    | 90H         | IA32_PERF_GLOBAL_STATUS |
| 38H         | R/EDI    | 98H         | PSDLA                   |
| 40H         | R/EBP    | A0H         | Reserved                |
| 48H         | R/ESP    | A8H         | Reserved                |
| 50H         | R8       | B0H         | EventingRIP             |
| 58H         | R9       | B8H         | Reserved                |

**18.4.1.1.2 Offcore Response Event**

Event number 0B7H support offcore response monitoring using an associated configuration MSR, MSR\_OFFCORE\_RSP0 (address 1A6H) in conjunction with umask value 01H or MSR\_OFFCORE\_RSP1 (address 1A7H) in conjunction with umask value 02H. Table 18-47 lists the event code, mask value and additional off-core configuration MSR that must be programmed to count off-core response events using IA32\_PMCx.

**Table 18-47. OffCore Response Event Encoding**

| Counter | Event code | UMask | Required Off-core Response MSR  |
|---------|------------|-------|---------------------------------|
| PMC0-1  | B7H        | 01H   | MSR_OFFCORE_RSP0 (address 1A6H) |
| PMC0-1  | B7H        | 02H   | MSR_OFFCORE_RSP1 (address 1A7H) |

Some of the MSR\_OFFCORE\_RESP [0,1] register bits are not valid in this processor and their use is reserved. The layout of MSR\_OFFCORE\_RSP0 and MSR\_OFFCORE\_RSP1 registers are defined in Table 18-48. Bits 15:0 specifies the request type of a transaction request to the uncore. Bits 30:16 specifies supplier information, bits 37:31 specifies snoop response information.

Additionally, MSR\_OFFCORE\_RSP0 provides bit 38 to enable measurement of average latency of specific type of offcore transaction requests using two programmable counter simultaneously, see Section 18.5.2.3 for details.

Table 18-48. Bit fields of the MSR\_OFFCORE\_RESP [0, 1] Registers

| Main          | Sub-field                         | Bit | Name                                | Description   |
|---------------|-----------------------------------|-----|-------------------------------------|---|
| Request Type  |                                   | 0   | DEMAND_DATA_RD                      | Demand cacheable data and L1 prefetch data reads.   |
|               |                                   | 1   | DEMAND_RFO                          | Demand cacheable data writes.   |
|               |                                   | 2   | DEMAND_CODE_RD                      | Demand code reads and prefetch code reads.  |
|               |                                   | 3   | Reserved                            | Reserved.   |
|               |                                   | 4   | Reserved                            | Reserved.   |
|               |                                   | 5   | PF_L2_RFO                           | L2 data RFO prefetches (includes PREFETCHW instruction).  |
|               |                                   | 6   | PF_L2_CODE_RD                       | L2 code HW prefetches.  |
|               |                                   | 7   | PARTIAL_READS                       | Partial reads (UC or WC).   |
|               |                                   | 8   | PARTIAL_WRITES                      | Partial writes (UC or WT or WP). Valid only for OFFCORE_RESP_1 event. Should only be used on PMC1. This bit is reserved for OFFCORE_RESP_0 event. |
|               |                                   | 9   | UC_CODE_READS                       | UC code reads.  |
|               |                                   | 10  | BUS_LOCKS                           | Bus locks and split lock requests.  |
|               |                                   | 11  | FULL_STREAMING_STORES               | Full streaming stores (WC). Valid only for OFFCORE_RESP_1 event. Should only be used on PMC1. This bit is reserved for OFFCORE_RESP_0 event.      |
|               |                                   | 12  | SW_PREFETCH                         | Software prefetches.  |
|               |                                   | 13  | PF_L1_DATA_RD                       | L1 data HW prefetches.  |
|               |                                   | 14  | PARTIAL_STREAMING_STORES            | Partial streaming stores (WC). Valid only for OFFCORE_RESP_1 event. Should only be used on PMC1. This bit is reserved for OFFCORE_RESP_0 event.   |
| Response Type | Any                               | 16  | ANY_RESPONSE                        | Account for any response.   |
|               | Data Supply from Untile           | 17  | NO_SUPP                             | No Supplier Details.  |
|               |                                   | 18  | Reserved                            | Reserved.   |
|               |                                   | 19  | L2_HIT_OTHER_TILE_NEAR              | Other tile L2 hit E Near.   |
|               |                                   | 20  | Reserved                            | Reserved.   |
|               |                                   | 21  | MCDRAM_NEAR                         | MCDRAM Local.   |
|               |                                   | 22  | MCDRAM_FAR_OR_L2_HIT_OTHER_TILE_FAR | MCDRAM Far or Other tile L2 hit far.  |
|               |                                   | 23  | DRAM_NEAR                           | DRAM Local.   |
|               |                                   | 24  | DRAM_FAR                            | DRAM Far.   |
|               | Data Supply from within same tile | 25  | L2_HITM_THIS_TILE                   | M-state.  |
|               |                                   | 26  | L2_HITE_THIS_TILE                   | E-state.  |
|               |                                   | 27  | L2_HITS_THIS_TILE                   | S-state.  |
|               |                                   | 28  | L2_HITF_THIS_TILE                   | F-state.  |
|               |                                   | 29  | Reserved                            | Reserved.   |
|               |                                   | 30  | Reserved                            | Reserved.   |

**Table 18-48. Bit fields of the MSR\_OFFCORE\_RESP [0, 1] Registers (Contd.)**

| Main                 | Sub-field   | Bit | Name   | Description   |
|----------------------|---|-----|--|---|
|                      | Snoop Info; Only Valid in case of Data Supply from Untile | 31  | SNOOP_NONE   | None of the cores were snooped.   |
|                      |   | 32  | NO_SNOOP_NEEDED  | No snoop was needed to satisfy the request.   |
|                      |   | 33  | Reserved   | Reserved.   |
|                      |   | 34  | Reserved   | Reserved.   |
|                      |   | 35  | HIT_OTHER_TILE_FWD   | Snoop request hit in the other tile with data forwarded.  |
|                      |   | 36  | HITM_OTHER_TILE  | A snoop was needed and it HitM-ed in other core's L1 cache. HitM denotes a cache-line was in modified state before effect as a result of snoop.   |
|                      |   | 37  | NON_DRAM   | Target was non-DRAM system address. This includes MMIO transactions.  |
| Outstanding requests | Weighted cycles   | 38  | OUTSTANDING (Valid only for MSR_OFFCORE_RESP0. Should only be used on PMCO. This bit is reserved for MSR_OFFCORE_RESP1). | If set, counts total number of weighted cycles of any outstanding offcore requests with data response. Valid only for OFFCORE_RESP_0 event. Should only be used on PMCO. This bit is reserved for OFFCORE_RESP_1 event. |

#### 18.4.1.1.3 Average Offcore Request Latency Measurement

Measurement of average latency of offcore transaction requests can be enabled using MSR\_OFFCORE\_RSP0.[bit 38] with the choice of request type specified in MSR\_OFFCORE\_RSP0.[bit 15:0].

Refer to Section 18.5.2.3, "Average Offcore Request Latency Measurement," for typical usage. Note that MSR\_OFFCORE\_RESPx registers are not shared between cores in Knights Landing. This allows one core to measure average latency while other core is measuring different offcore response events.

## 18.5 PERFORMANCE MONITORING (INTEL® ATOM™ PROCESSORS)

### 18.5.1 Performance Monitoring (45 nm and 32 nm Intel® Atom™ Processors)

45 nm and 32 nm Intel Atom processors report architectural performance monitoring versionID = 3 (supporting the aggregate capabilities of versionID 1, 2, and 3; see Section 18.2.3) and a host of non-architectural monitoring capabilities. These 45 nm and 32 nm Intel Atom processors provide two general-purpose performance counters (IA32\_PMC0, IA32\_PMC1) and three fixed-function performance counters (IA32\_FIXED\_CTR0, IA32\_FIXED\_CTR1, IA32\_FIXED\_CTR2).

Non-architectural performance monitoring in Intel Atom processor family uses the IA32\_PERFEVTSELx MSR to configure a set of non-architecture performance monitoring events to be counted by the corresponding general-purpose performance counter. The list of non-architectural performance monitoring events is listed in Table 19-29.

Architectural and non-architectural performance monitoring events in 45 nm and 32 nm Intel Atom processors support thread qualification using bit 21 (AnyThread) of IA32\_PERFEVTSELx MSR, i.e. if IA32\_PERFEVTSELx.AnyThread = 1, event counts include monitored conditions due to either logical processors in the same processor core.

The bit fields within each IA32\_PERFEVTSELx MSR are defined in Figure 18-6 and described in Section 18.2.1.1 and Section 18.2.3.

Valid event mask (Umask) bits are listed in Chapter 19. The UMASK field may contain sub-fields that provide the same qualifying actions like those listed in Table 18-61, Table 18-62, Table 18-63, and Table 18-64. One or more of these sub-fields may apply to specific events on an event-by-event basis. Details are listed in Table 19-29 in

Chapter 19, “Performance Monitoring Events.” Precise Event Based Monitoring is supported using IA32\_PMC0 (see also Section 17.4.9, “BTS and DS Save Area”).

## 18.5.2 Performance Monitoring for Silvermont Microarchitecture

Intel processors based on the Silvermont microarchitecture report architectural performance monitoring versionID = 3 (see Section 18.2.3) and a host of non-architectural monitoring capabilities. Intel processors based on the Silvermont microarchitecture provide two general-purpose performance counters (IA32\_PMC0, IA32\_PMC1) and three fixed-function performance counters (IA32\_FIXED\_CTR0, IA32\_FIXED\_CTR1, IA32\_FIXED\_CTR2). Intel Atom processors based on the Airmont microarchitecture support the same performance monitoring capabilities as those based on the Silvermont microarchitecture.

Non-architectural performance monitoring in the Silvermont microarchitecture uses the IA32\_PERFEVTSELx MSR to configure a set of non-architecture performance monitoring events to be counted by the corresponding general-purpose performance counter. The list of non-architectural performance monitoring events is listed in Table 19-28.

The bit fields (except bit 21) within each IA32\_PERFEVTSELx MSR are defined in Figure 18-6 and described in Section 18.2.1.1 and Section 18.2.3. Architectural and non-architectural performance monitoring events in the Silvermont microarchitecture ignore the AnyThread qualification regardless of its setting in IA32\_PERFEVTSELx MSR.

### 18.5.2.1 Enhancements of Performance Monitoring in the Processor Core

The notable enhancements in the monitoring of performance events in the processor core include:

- The width of counter reported by CPUID.0AH:EAX[23:16] is 40 bits.
- Off-core response counting facility. This facility in the processor core allows software to count certain transaction responses between the processor core to sub-systems outside the processor core (uncore). Counting off-core response requires additional event qualification configuration facility in conjunction with IA32\_PERFEVTSELx. Two off-core response MSRs are provided to use in conjunction with specific event codes that must be specified with IA32\_PERFEVTSELx.
- Average request latency measurement. The off-core response counting facility can be combined to use two performance counters to count the occurrences and weighted cycles of transaction requests.

#### 18.5.2.1.1 Processor Event Based Sampling (PEBS)

In the Silvermont microarchitecture, the PEBS facility can be used with precise events. PEBS is supported using IA32\_PMC0 (see also Section 17.4.9).

PEBS uses a debug store mechanism to store a set of architectural state information for the processor. The information provides architectural state of the instruction executed after the instruction that caused the event (See Section 18.6.2.4).

The list of precise events supported in the Silvermont microarchitecture is shown in Table 18-49.

**Table 18-49. PEBS Performance Events for the Silvermont Microarchitecture**

| Event Name      | Event Select | Sub-event      | UMask |
|-----------------|--------------|----------------|-------|
| BR_INST_RETIRED | C4H          | ALL_BRANCHES   | 00H   |
|                 |              | JCC            | 7EH   |
|                 |              | TAKEN_JCC      | FEH   |
|                 |              | CALL           | F9H   |
|                 |              | REL_CALL       | FDH   |
|                 |              | IND_CALL       | FBH   |
|                 |              | NON_RETURN_IND | EBH   |
|                 |              | FAR_BRANCH     | BFH   |
|                 |              | RETURN         | F7H   |

**Table 18-49. PEBS Performance Events for the Silvermont Microarchitecture (Contd.)**

| Event Name       | Event Select | Sub-event           | UMask |
|------------------|--------------|---------------------|-------|
| BR_MISP_RETIRED  | C5H          | ALL_BRANCHES        | 00H   |
|                  |              | JCC                 | 7EH   |
|                  |              | TAKEN_JCC           | FEH   |
|                  |              | IND_CALL            | FBH   |
|                  |              | NON_RETURN_IND      | EBH   |
|                  |              | RETURN              | F7H   |
| MEM_UOPS_RETIRED | 04H          | L2_HIT_LOADS        | 02H   |
|                  |              | L2_MISS_LOADS       | 04H   |
|                  |              | DLTB_MISS_LOADS     | 08H   |
|                  |              | HITM                | 20H   |
| REHABQ           | 03H          | LD_BLOCK_ST_FORWARD | 01H   |
|                  |              | LD_SPLITS           | 08H   |

PEBS Record Format The PEBS record format supported by processors based on the Intel Silvermont microarchitecture is shown in Table 18-50, and each field in the PEBS record is 64 bits long.

**Table 18-50. PEBS Record Format for the Silvermont Microarchitecture**

| Byte Offset | Field    | Byte Offset | Field                   |
|-------------|----------|-------------|-------------------------|
| 00H         | R/EFLAGS | 60H         | R10                     |
| 08H         | R/EIP    | 68H         | R11                     |
| 10H         | R/EAX    | 70H         | R12                     |
| 18H         | R/EBX    | 78H         | R13                     |
| 20H         | R/ECX    | 80H         | R14                     |
| 28H         | R/EDX    | 88H         | R15                     |
| 30H         | R/ESI    | 90H         | IA32_PERF_GLOBAL_STATUS |
| 38H         | R/EDI    | 98H         | Reserved                |
| 40H         | R/EBP    | A0H         | Reserved                |
| 48H         | R/ESP    | A8H         | Reserved                |
| 50H         | R8       | B0H         | EventingRIP             |
| 58H         | R9       | B8H         | Reserved                |

### 18.5.2.2 Offcore Response Event

Event number 0B7H support offcore response monitoring using an associated configuration MSR, MSR\_OFFCORE\_RSP0 (address 1A6H) in conjunction with umask value 01H or MSR\_OFFCORE\_RSP1 (address 1A7H) in conjunction with umask value 02H. Table 18-51 lists the event code, mask value and additional off-core configuration MSR that must be programmed to count off-core response events using IA32\_PMCx.




In the Silvermont microarchitecture, each MSR\_OFFCORE\_RSPx is shared by two processor cores.

**Table 18-51. OffCore Response Event Encoding**

| Counter | Event code | UMask | Required Off-core Response MSR  |
|---------|------------|-------|---------------------------------|
| PMCO-1  | B7H        | 01H   | MSR_OFFCORE_RSP0 (address 1A6H) |
| PMCO-1  | B7H        | 02H   | MSR_OFFCORE_RSP1 (address 1A7H) |

The layout of MSR\_OFFCORE\_RSP0 and MSR\_OFFCORE\_RSP1 are shown in Figure 18-37 and Figure 18-38. Bits 15:0 specifies the request type of a transaction request to the uncore. Bits 30:16 specifies supplier information, bits 37:31 specifies snoop response information.

Additionally, MSR\_OFFCORE\_RSP0 provides bit 38 to enable measurement of average latency of specific type of offcore transaction requests using two programmable counter simultaneously, see Section 18.5.2.3 for details.



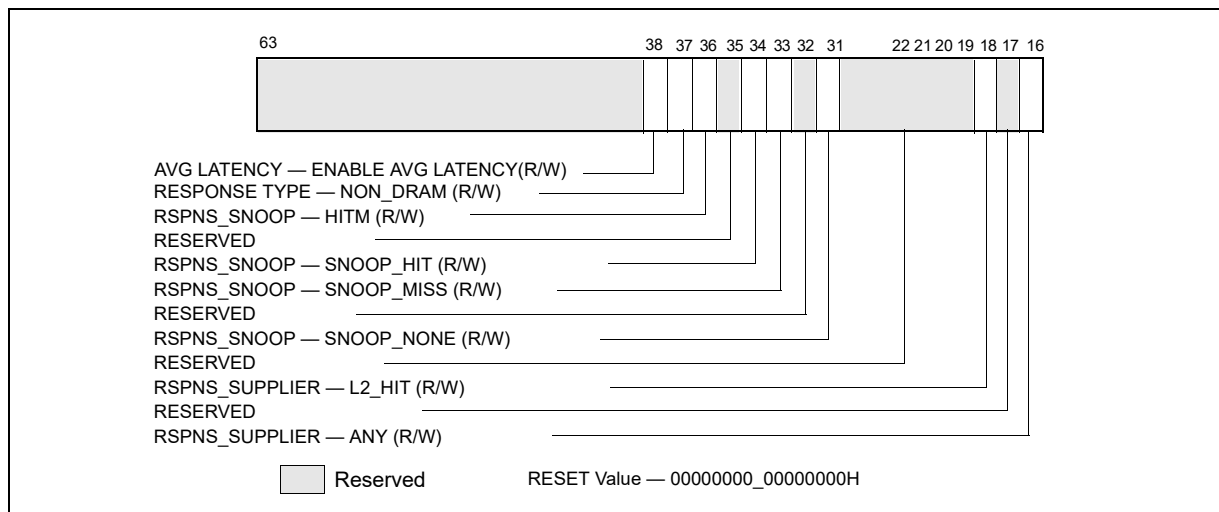
**Figure 18-37. Request\_Type Fields for MSR\_OFFCORE\_RSPx**

**Table 18-52. MSR\_OFFCORE\_RSPx Request\_Type Field Definition**

| Bit Name     | Offset | Description   |
|--------------|--------|---|
| DMND_DATA_RD | 0      | (R/W). Counts the number of demand and DCU prefetch data reads of full and partial cachelines as well as demand data page table entry cacheline reads. Does not count L2 data read prefetches or instruction fetches. |
| DMND_RFO     | 1      | (R/W). Counts the number of demand and DCU prefetch reads for ownership (RFO) requests generated by a write to data cacheline. Does not count L2 RFO prefetches.  |
| DMND_IFETCH  | 2      | (R/W). Counts the number of demand and DCU prefetch instruction cacheline reads. Does not count L2 code read prefetches.  |
| WB           | 3      | (R/W). Counts the number of writeback (modified to exclusive) transactions.   |
| PF_DATA_RD   | 4      | (R/W). Counts the number of data cacheline reads generated by L2 prefetchers.   |
| PF_RFO       | 5      | (R/W). Counts the number of RFO requests generated by L2 prefetchers.   |
| PF_IFETCH    | 6      | (R/W). Counts the number of code reads generated by L2 prefetchers.   |
| PARTIAL_READ | 7      | (R/W). Counts the number of demand reads of partial cache lines (including UC and WC).  |

**Table 18-52. MSR\_OFFCORE\_RSPx Request\_Type Field Definition (Contd.)**

| Bit Name        | Offset | Description  |
|-----------------|--------|--|
| PARTIAL_WRITE   | 8      | (R/W). Counts the number of demand RFO requests to write to partial cache lines (includes UC, WT and WP) |
| UC_IFETCH       | 9      | (R/W). Counts the number of UC instruction fetches.  |
| BUS_LOCKS       | 10     | (R/W). Bus lock and split lock requests  |
| STRM_ST         | 11     | (R/W). Streaming store requests  |
| SW_PREFETCH     | 12     | (R/W). Counts software prefetch requests   |
| PF_DATA_RD      | 13     | (R/W). Counts DCU hardware prefetcher data read requests   |
| PARTIAL_STRM_ST | 14     | (R/W). Streaming store requests  |
| ANY             | 15     | (R/W). Any request that crosses IDI, including I/O.  |



**Figure 18-38. Response\_Supplier and Snoop Info Fields for MSR\_OFFCORE\_RSPx**

To properly program this extra register, software must set at least one request type bit (Table 18-52) and a valid response type pattern (Table 18-53, Table 18-54). Otherwise, the event count reported will be zero. It is permissible and useful to set multiple request and response type bits in order to obtain various classes of off-core response events. Although MSR\_OFFCORE\_RSPx allow an agent software to program numerous combinations that meet the above guideline, not all combinations produce meaningful data.

**Table 18-53. MSR\_OFFCORE\_RSP\_x Response Supplier Info Field Definition**

| Subtype       | Bit Name     | Offset | Description   |
|---------------|--------------|--------|---|
| Common        | ANY_RESPONSE | 16     | (R/W). Catch all value for any response types.        |
| Supplier Info | Reserved     | 17     | Reserved  |
|               | L2_HIT       | 18     | (R/W). Cache reference hit L2 in either M/E/S states. |
|               | Reserved     | 30:19  | Reserved  |

To specify a complete offcore response filter, software must properly program bits in the request and response type fields. A valid request type must have at least one bit set in the non-reserved bits of 15:0. A valid response type must be a non-zero value of the following expression:

ANY | [(‘OR’ of Supplier Info Bits) & (‘OR’ of Snoop Info Bits)]

If “ANY” bit is set, the supplier and snoop info bits are ignored.

**Table 18-54. MSR\_OFFCORE\_RSPx Snoop Info Field Definition**

| Subtype    | Bit Name    | Offset | Description  |
|------------|-------------|--------|--|
| Snoop Info | SNP_NONE    | 31     | (R/W). No details on snoop-related information.  |
|            | Reserved    | 32     | Reserved   |
|            | SNOOP_MISS  | 33     | (R/W). Counts the number of snoop misses when L2 misses.   |
|            | SNOOP_HIT   | 34     | (R/W). Counts the number of snoops hit in the other module where no modified copies were found.  |
|            | Reserved    | 35     | Reserved   |
|            | HITM        | 36     | (R/W). Counts the number of snoops hit in the other module where modified copies were found in other core's L1 cache.  |
|            | NON_DRAM    | 37     | (R/W). Target was non-DRAM system address. This includes MMIO transactions.  |
|            | AVG_LATENCY | 38     | (R/W). Enable average latency measurement by counting weighted cycles of outstanding offcore requests of the request type specified in bits 15:0 and any response (bits 37:16 cleared to 0).<br><br>This bit is available in MSR_OFFCORE_RESP0. The weighted cycles is accumulated in the specified programmable counter IA32_PMCx and the occurrence of specified requests are counted in the other programmable counter. |

### 18.5.2.3 Average Offcore Request Latency Measurement

Average latency for offcore transactions can be determined by using both MSR\_OFFCORE\_RSP registers. Using two performance monitoring counters, program the two OFFCORE\_RESPONSE event encodings into the corresponding IA32\_PERFVTSELx MSRs. Count the weighted cycles via MSR\_OFFCORE\_RSP0 by programming a request type in MSR\_OFFCORE\_RSP0.[15:0] and setting MSR\_OFFCORE\_RSP0.OUTSTANDING[38] to 1, while setting the remaining bits to 0. Count the number of requests via MSR\_OFFCORE\_RSP1 by programming the same request type from MSR\_OFFCORE\_RSP0 into MSR\_OFFCORE\_RSP1[bit 15:0], and setting MSR\_OFFCORE\_RSP1.ANY\_RESPONSE[16] = 1, while setting the remaining bits to 0. The average latency can be obtained by dividing the value of the IA32\_PMCx register that counted weight cycles by the register that counted requests.

## 18.5.3 Performance Monitoring for Goldmont Microarchitecture

Intel Atom processors based on the Goldmont microarchitecture report architectural performance monitoring versionID = 4 (see Section 18.2.4) and support non-architectural monitoring capabilities described in this section.

Architectural performance monitoring version 4 capabilities are described in Section 18.2.4.

The bit fields (except bit 21) within each IA32\_PERFVTSELx MSR are defined in Figure 18-6 and described in Section 18.2.1.1 and Section 18.2.3. The Goldmont microarchitecture does not support Hyper-Threading and thus architectural and non-architectural performance monitoring events ignore the AnyThread qualification regardless of its setting in the IA32\_PERFVTSELx MSR. However, Goldmont does not set the AnyThread deprecation bit (CPUID.0AH:EDX[15]).

The core PMU's capability is similar to that of the Silvermont microarchitecture described in Section 18.5.2, with some differences and enhancements summarized in Table 18-55.

**Table 18-55. Core PMU Comparison Between the Goldmont and Silvermont Microarchitectures**

| Box  | The Goldmont microarchitecture  | The Silvermont microarchitecture  | Comment  |
|--|---|---|--|
| # of Fixed counters per core   | 3   | 3   | Use CPUID to enumerate # of counters.  |
| # of general-purpose counters per core                               | 4   | 2   |  |
| Counter width (R,W)  | R:48, W: 32/48  | R:40, W:32  | See Section 18.2.2.  |
| Architectural Performance Monitoring version ID                      | 4   | 3   | Use CPUID to enumerate # of counters.  |
| PMI Overhead Mitigation  | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with streamlined semantics.</li> <li>▪ Freeze_LBR_on_PMI with streamlined semantics for branch profiling.</li> </ul>     | <ul style="list-style-type: none"> <li>▪ Freeze_Perfmon_on_PMI with legacy semantics.</li> <li>▪ Freeze_LBR_on_PMI with legacy semantics for branch profiling.</li> </ul> | See Section 17.4.7. Legacy semantics not supported with version 4 or higher. |
| Counter and Buffer Overflow Status Management                        | <ul style="list-style-type: none"> <li>▪ Query via IA32_PERF_GLOBAL_STATUS</li> <li>▪ Reset via IA32_PERF_GLOBAL_STATUS_RESET</li> <li>▪ Set via IA32_PERF_GLOBAL_STATUS_SET</li> </ul> | <ul style="list-style-type: none"> <li>▪ Query via IA32_PERF_GLOBAL_STATUS</li> <li>▪ Reset via IA32_PERF_GLOBAL_OVF_CTRL</li> </ul>                                      | See Section 18.2.4.  |
| IA32_PERF_GLOBAL_STATUS Indicators of Overflow/Overhead/Interference | <ul style="list-style-type: none"> <li>▪ Individual counter overflow</li> <li>▪ PEBS buffer overflow</li> <li>▪ ToPA buffer overflow</li> <li>▪ CTR_Frz, LBR_Frz</li> </ul>             | <ul style="list-style-type: none"> <li>▪ Individual counter overflow</li> <li>▪ PEBS buffer overflow</li> </ul>   | See Section 18.2.4.  |
| Enable control in IA32_PERF_GLOBAL_STATUS                            | <ul style="list-style-type: none"> <li>▪ CTR_Frz,</li> <li>▪ LBR_Frz</li> </ul>   | No  | See Section 18.2.4.1.  |
| Perfmon Counter In-Use Indicator                                     | Query IA32_PERF_GLOBAL_INUSE  | No  | See Section 18.2.4.3.  |
| Processor Event Based Sampling (PEBS) Events                         | General-Purpose Counter 0 only. Supports all events (precise and non-precise). Precise events are listed in Table 18-56.  | See Section 18.5.2.1.1. General-Purpose Counter 0 only. Only supports precise events (see Table 18-49).   | IA32_PMC0 only.  |
| PEBS record format encoding  | 0011b   | 0010b   |  |
| Reduce skid PEBS   | IA32_PMC0 only  | No  |  |
| Data Address Profiling   | Yes   | No  |  |
| PEBS record layout   | Table 18-57; enhanced fields at offsets 90H- 98H; and TSC record field at C0H.  | Table 18-50.  |  |
| PEBS EventingIP  | Yes   | Yes   |  |
| Off-core Response Event  | MSR 1A6H and 1A7H, each core has its own register.  | MSR 1A6H and 1A7H, shared by a pair of cores.   | Nehalem supports 1A6H only.  |

### 18.5.3.1 Processor Event Based Sampling (PEBS)

Processor event based sampling (PEBS) on the Goldmont microarchitecture is enhanced over prior generations with respect to sampling support of precise events and non-precise events. In the Goldmont microarchitecture, PEBS is supported using IA32\_PMC0 for all events (see Section 17.4.9).

PEBS uses a debug store mechanism to store a set of architectural state information for the processor at the time the sample was generated.

Precise events work the same way on Goldmont microarchitecture as on the Silvermont microarchitecture. The record will be generated after an instruction that causes the event when the counter is already overflowed and will capture the architectural state at this point (see Section 18.6.2.4 and Section 17.4.9). The eventingIP in the record will indicate the instruction that caused the event. The list of precise events supported in the Goldmont microarchitecture is shown in Table 18-56.

In the Goldmont microarchitecture, the PEBS facility also supports the use of non-precise events to record processor state information into PEBS records with the same format as with precise events.

However, a non-precise event may not be attributable to a particular retired instruction or the time of instruction execution. When the counter overflows, a PEBS record will be generated at the next opportunity. Consider the event ICACHE.HIT. When the counter overflows, the processor is fetching future instructions. The PEBS record will be generated at the next opportunity and capture the state at the processor's current retirement point. It is likely that the instruction fetch that caused the event to increment was beyond that current retirement point. Other examples of non-precise events are CPU\_CLK\_UNHALTED.CORE\_P and HARDWARE\_INTERRUPTS.RECEIVED. CPU\_CLK\_UNHALTED.CORE\_P will increment each cycle that the processor is awake. When the counter overflows, there may be many instructions in various stages of execution. Additionally, zero, one or multiple instructions may be retired the cycle that the counter overflows. HARDWARE\_INTERRUPTS.RECEIVED increments independent of any instructions being executed. For all non-precise events, the PEBS record will be generated at the next opportunity, after the counter has overflowed. The PEBS facility thus allows for identification of the instructions which were executing when the event overflowed.

After generating a record for a non-precise event, the PEBS facility reloads the counter and resumes execution, just as is done for precise events. Unlike interrupt-based sampling, which requires an interrupt service routine to collect the sample and reload the counter, the PEBS facility can collect samples even when interrupts are masked and without using NMI. Since a PEBS record is generated immediately when a counter for a non-precise event is enabled, it may also be generated after an overflow is set by an MSR write to IA32\_PERF\_GLOBAL\_STATUS\_SET.

**Table 18-56. Precise Events Supported by the Goldmont Microarchitecture**

| Event Name       | Event Select | Sub-event        | UMask |
|------------------|--------------|------------------|-------|
| LD_BLOCKS        | 03H          | DATA_UNKNOWN     | 01H   |
|                  |              | STORE_FORWARD    | 02H   |
|                  |              | 4K_ALIAS         | 04H   |
|                  |              | UTLB_MISS        | 08H   |
|                  |              | ALL_BLOCK        | 10H   |
| MISALIGN_MEM_REF | 13H          | LOAD_PAGE_SPLIT  | 02H   |
|                  |              | STORE_PAGE_SPLIT | 04H   |
| INST_RETIRED     | COH          | ANY              | 00H   |
| UOPS_RETITRED    | C2H          | ANY              | 00H   |
|                  |              | LD_SPLITSMS      | 01H   |
| BR_INST_RETIRED  | C4H          | ALL_BRANCHES     | 00H   |
|                  |              | JCC              | 7EH   |
|                  |              | TAKEN_JCC        | FEH   |
|                  |              | CALL             | F9H   |
|                  |              | REL_CALL         | FDH   |
|                  |              | IND_CALL         | FBH   |

**Table 18-56. Precise Events Supported by the Goldmont Microarchitecture (Contd.)**

| Event Name            | Event Select | Sub-event        | UMask |
|-----------------------|--------------|------------------|-------|
|                       |              | NON_RETURN_IND   | EBH   |
|                       |              | FAR_BRANCH       | BFH   |
|                       |              | RETURN           | F7H   |
| BR_MISP_RETIRED       | C5H          | ALL_BRANCHES     | 00H   |
|                       |              | JCC              | 7EH   |
|                       |              | TAKEN_JCC        | FEH   |
|                       |              | IND_CALL         | FBH   |
|                       |              | NON_RETURN_IND   | EBH   |
|                       |              | RETURN           | F7H   |
| MEM_UOPS_RETIRED      | D0H          | ALL_LOADS        | 81H   |
|                       |              | ALL_STORES       | 82H   |
|                       |              | ALL              | 83H   |
|                       |              | DLTB_MISS_LOADS  | 11H   |
|                       |              | DLTB_MISS_STORES | 12H   |
|                       |              | DLTB_MISS        | 13H   |
| MEM_LOAD_UOPS_RETIRED | D1H          | L1_HIT           | 01H   |
|                       |              | L2_HIT           | 02H   |
|                       |              | L1_MISS          | 08H   |
|                       |              | L2_MISS          | 10H   |
|                       |              | HITM             | 20H   |
|                       |              | WCB_HIT          | 40H   |
|                       |              | DRAM_HIT         | 80H   |

The PEBS record format supported by processors based on the Intel Goldmont microarchitecture is shown in Table 18-57, and each field in the PEBS record is 64 bits long.

**Table 18-57. PEBS Record Format for the Goldmont Microarchitecture**

| Byte Offset | Field    | Byte Offset | Field               |
|-------------|----------|-------------|---------------------|
| 00H         | R/EFLAGS | 68H         | R11                 |
| 08H         | R/EIP    | 70H         | R12                 |
| 10H         | R/EAX    | 78H         | R13                 |
| 18H         | R/EBX    | 80H         | R14                 |
| 20H         | R/ECX    | 88H         | R15                 |
| 28H         | R/EDX    | 90H         | Applicable Counters |
| 30H         | R/ESI    | 98H         | Data Linear Address |
| 38H         | R/EDI    | A0H         | Reserved            |
| 40H         | R/EBP    | A8H         | Reserved            |
| 48H         | R/ESP    | B0H         | EventingRIP         |
| 50H         | R8       | B8H         | Reserved            |
| 58H         | R9       | C0H         | TSC                 |
| 60H         | R10      |             |                     |

On Goldmont microarchitecture, all 64 bits of architectural registers are written into the PEBS record regardless of processor mode.

With PEBS record format encoding 0011b, offset 90H reports the "Applicable Counter" field, which indicates which counters actually requested generating a PEBS record. This allows software to correlate the PEBS record entry properly with the instruction that caused the event even when multiple counters are configured to record PEBS records and multiple bits are set in the field. Additionally, offset C0H captures a snapshot of the TSC that provides a time line annotation for each PEBS record entry.

#### 18.5.3.1.1 PEBS Data Linear Address Profiling

Goldmont supports the Data Linear Address field introduced in Haswell. It does not support the Data Source Encoding or Latency Value fields that are also part of Data Address Profiling; those fields are present in the record but are reserved.

For Goldmont microarchitecture, the Data Linear Address field will record the linear address of memory accesses in the previous instruction (e.g. the one that triggered a precise event that caused the PEBS record to be generated). Goldmont microarchitecture may record a Data Linear Address for the instruction that caused the event even for events not related to memory accesses. This may differ from other microarchitectures.

#### 18.5.3.1.2 Reduced Skid PEBS

For precise events, upon triggering a PEBS assist, there will be a finite delay between the time the counter overflows and when the microcode starts to carry out its data collection obligations. The Reduced Skid mechanism mitigates the "skid" problem by providing an early indication of when the counter is about to overflow, allowing the machine to more precisely trap on the instruction that actually caused the counter overflow thus greatly reducing skid.

This mechanism is a superset of the PDIR mechanism available in the Sandy Bridge microarchitecture. See Section 18.3.4.4.4

In the Goldmont microarchitecture, the mechanism applies to all precise events including, INST\_RETIRE, except for UOPS\_RETIRE. However, the Reduced Skid mechanism is disabled for any counter when the INV, ANY, E, or CMASK fields are set.

For the Reduced Skid mechanism to operate correctly, the performance monitoring counters should not be reconfigured or modified when they are running with PEBS enabled. The counters need to be disabled (e.g. via IA32\_PERF\_GLOBAL\_CTRL MSR) before changes to the configuration (e.g. what event is specified in IA32\_PERFEVTSELx or whether PEBS is enabled for that counter via IA32\_PEBS\_ENABLE) or counter value (MSR write to IA32\_PMCx and IA32\_A\_PMCx).

#### 18.5.3.1.3 Enhancements to IA32\_PERF\_GLOBAL\_STATUS.OvfDSBuffer[62]

In addition to IA32\_PERF\_GLOBAL\_STATUS.OvfDSBuffer[62] being set when PEBS\_Index reaches the PEBS\_Interrupt\_Threshold, the bit is also set when PEBS\_Index is out of bounds. That is, the bit will be set when PEBS\_Index < PEBS\_Buffer\_Base or PEBS\_Index > PEBS\_Absolute\_Maximum. Note that when an out of bound condition is encountered, the overflow bits in IA32\_PERF\_GLOBAL\_STATUS will be cleared according to Applicable Counters, however the IA32\_PMCx values will not be reloaded with the Reset values stored in the DS\_AREA.

### 18.5.3.2 Offcore Response Event

Event number 0B7H support offcore response monitoring using an associated configuration MSR, MSR\_OFFCORE\_RSP0 (address 1A6H) in conjunction with umask value 01H or MSR\_OFFCORE\_RSP1 (address 1A7H) in conjunction with umask value 02H. Table 18-51 lists the event code, mask value and additional off-core configuration MSR that must be programmed to count off-core response events using IA32\_PMCx.

The Goldmont microarchitecture provides unique pairs of MSR\_OFFCORE\_RSPx registers per core.

The layout of MSR\_OFFCORE\_RSP0 and MSR\_OFFCORE\_RSP1 are organized as follows:

- Bits 15:0 specifies the request type of a transaction request to the uncore. This is described in Table 18-58.
- Bits 30:16 specifies common supplier information or an L2 Hit, and is described in Table 18-53.

- If L2 misses, then Bits 37:31 can be used to specify snoop response information and is described in Table 18-59.
- For outstanding requests, bit 38 can enable measurement of average latency of specific type of offcore transaction requests using two programmable counter simultaneously; see Section 18.5.2.3 for details.

**Table 18-58. MSR\_OFFCORE\_RSPx Request\_Type Field Definition**

| Bit Name                 | Offset | Description  |
|--------------------------|--------|--|
| DEMAND_DATA_RD           | 0      | (R/W) Counts cacheline read requests due to demand reads (excludes prefetches).  |
| DEMAND_RFO               | 1      | (R/W) Counts cacheline read for ownership (RFO) requests due to demand writes (excludes prefetches).                         |
| DEMAND_CODE_RD           | 2      | (R/W) Counts demand instruction cacheline and I-side prefetch requests that miss the instruction cache.                      |
| COREWB                   | 3      | (R/W) Counts writeback transactions caused by L1 or L2 cache evictions.  |
| PF_L2_DATA_RD            | 4      | (R/W) Counts data cacheline reads generated by hardware L2 cache prefetcher.   |
| PF_L2_RFO                | 5      | (R/W) Counts reads for ownership (RFO) requests generated by L2 prefetcher.  |
| Reserved                 | 6      | Reserved.  |
| PARTIAL_READS            | 7      | (R/W) Counts demand data partial reads, including data in uncacheable (UC) or uncacheable (WC) write combining memory types. |
| PARTIAL_WRITES           | 8      | (R/W) Counts partial writes, including uncacheable (UC), write through (WT) and write protected (WP) memory type writes.     |
| UC_CODE_READS            | 9      | (R/W) Counts code reads in uncacheable (UC) memory region.   |
| BUS_LOCKS                | 10     | (R/W) Counts bus lock and split lock requests.   |
| FULL_STREAMING_STORES    | 11     | (R/W) Counts full cacheline writes due to streaming stores.  |
| SW_PREFETCH              | 12     | (R/W) Counts cacheline requests due to software prefetch instructions.   |
| PF_L1_DATA_RD            | 13     | (R/W) Counts data cacheline reads generated by hardware L1 data cache prefetcher.  |
| PARTIAL_STREAMING_STORES | 14     | (R/W) Counts partial cacheline writes due to streaming stores.   |
| ANY_REQUEST              | 15     | (R/W) Counts requests to the uncore subsystem.   |

To properly program this extra register, software must set at least one request type bit (Table 18-52) and a valid response type pattern (either Table 18-53 or Table 18-59). Otherwise, the event count reported will be zero. It is permissible and useful to set multiple request and response type bits in order to obtain various classes of off-core response events. Although MSR\_OFFCORE\_RSPx allow an agent software to program numerous combinations that meet the above guideline, not all combinations produce meaningful data.

**Table 18-59. MSR\_OFFCORE\_RSPx For L2 Miss and Outstanding Requests**

| Subtype                 | Bit Name                                  | Offset | Description   |
|-------------------------|---|--------|---|
| L2_MISS<br>(Snoop Info) | Reserved                                  | 32:31  | Reserved  |
|                         | L2_MISS.SNOOP_MISS_0<br>R_NO_SNOOP_NEEDED | 33     | (R/W) A true miss to this module, for which a snoop request missed the other module or no snoop was performed/needed. |
|                         | L2_MISS.HIT_OTHER_CO<br>RE_NO_FWD         | 34     | (R/W) A snoop hit in the other processor module, but no data forwarding is required.                                  |
|                         | Reserved                                  | 35     | Reserved  |
|                         | L2_MISS.HITM_OTHER_C<br>ORE               | 36     | (R/W) Counts the number of snoops hit in the other module or other core's L1 where modified copies were found.        |
|                         | L2_MISS.NON_DRAM                          | 37     | (R/W) Target was a non-DRAM system address. This includes MMIO transactions.  |



**Table 18-59. MSR\_OFFCORE\_RSPx For L2 Miss and Outstanding Requests (Contd.)**

| Subtype                           | Bit Name    | Offset | Description  |
|-----------------------------------|-------------|--------|--|
| Outstanding requests <sup>1</sup> | OUTSTANDING | 38     | (R/W) Counts weighted cycles of outstanding offcore requests of the request type specified in bits 15:0, from the time the XQ receives the request and any response is received. Bits 37:16 must be set to 0. This bit is only available in MSR_OFFCORE_RESPO. |

**NOTES:**

1. See Section 18.5.2.3, "Average Offcore Request Latency Measurement" for details on how to use this bit to extract average latency.

To specify a complete offcore response filter, software must properly program bits in the request and response type fields. A valid request type must have at least one bit set in the non-reserved bits of 15:0. A valid response type must be a non-zero value of the following expression:

[ANY 'OR' (L2 Hit) ] 'XOR' ( Snoop Info Bits) 'XOR' (Avg Latency)

**18.5.3.3 Average Offcore Request Latency Measurement**

In Goldmont microarchitecture, measurement of average latency of offcore transaction requests is the same as described in Section 18.5.2.3.

**18.5.4 Performance Monitoring for Goldmont Plus Microarchitecture**

Intel Atom processors based on the Goldmont Plus microarchitecture report architectural performance monitoring versionID = 4 and support non-architectural monitoring capabilities described in this section.

Architectural performance monitoring version 4 capabilities are described in Section 18.2.4.

Goldmont Plus performance monitoring capabilities are similar to Goldmont capabilities. The differences are in specific events and in which counters support PEBS. Goldmont Plus introduces the ability for fixed performance monitoring counters to generate PEBS records.

Goldmont Plus will set the AnyThread deprecation CPUID bit (CPUID.0AH:EDX[15]) to indicate that the Any-Thread bits in IA32\_PERFVTSELx and IA32\_FIXED\_CTR\_CTRL have no effect.

The core PMU's capability is similar to that of the Goldmont microarchitecture described in Section 18.6.3, with some differences and enhancements summarized in Table 18-60.

**Table 18-60. Core PMU Comparison Between the Goldmont Plus and Goldmont Microarchitectures**

| Box   | Goldmont Plus Microarchitecture   | Goldmont Microarchitecture   | Comment                                      |
|---|---|--|--|
| # of Fixed counters per core                    | 3   | 3  | No change.                                   |
| # of general-purpose counters per core          | 4   | 4  | No change.                                   |
| Counter width (R,W)                             | R:48, W: 32/48  | R:48, W: 32/48   | No change.                                   |
| Architectural Performance Monitoring version ID | 4   | 4  | No change.                                   |
| Processor Event Based Sampling (PEBS) Events    | All General-Purpose and Fixed counters. Each General-Purpose counter supports all events (precise and non-precise). | General-Purpose Counter 0 only. Supports all events (precise and non-precise). Precise events are listed in Table 18-56. | Goldmont Plus supports PEBS on all counters. |
| PEBS record format encoding                     | 0011b   | 0011b  | No change.                                   |

### 18.5.4.1 Extended PEBS

The Extended PEBS feature, introduced in Goldmont Plus microarchitecture, supports PEBS (Processor Event Based Sampling) on a fixed-function performance counters as well as all four general purpose counters (PMC0-3). PEBS can be enabled for the four general purpose counters using PEBS\_EN\_PMCi bits of IA32\_PEBS\_ENABLE (i = 0, 1, 2, 3). PEBS can be enabled for the 3 fixed function counters using the PEBS\_EN\_FIXEDi bits of IA32\_PEBS\_ENABLE (I = 0, 1, 2).




Figure 18-39. Layout of IA32\_PEBS\_ENABLE MSR

Similar to Goldmont microarchitecture, Goldmont Plus microarchitecture processors can generate PEBS record events on both precise as well as non-precise events.

A PEBS record due to a precise event will be generated after an instruction that causes the event when the counter has already overflowed. A PEBS record due to a non-precise event will occur at the next opportunity after the counter has overflowed, including immediately after an overflow is set by an MSR write.

IA32\_FIXED\_CTR0 counts instructions retired and is a precise event. IA32\_FIXED\_CTR1 counts unhalting core cycles and is a non-precise event. IA32\_FIXED\_CTR2 counts unhalting reference cycles and is a non-precise event.

The Applicable Counter field at offset 90H of the PEBS record indicates which counters caused the PEBS record to be generated. It is in the same format as the enable bits for each counter in IA32\_PEBS\_ENABLE. As an example, an Applicable Counter field with bits 2 and 32 set would indicate that both general purpose counter 2 and fixed function counter 0 generated the PEBS record.

- To properly use PEBS for the additional counters, software will need to set up the counter reset values in PEBS portion of the DS\_BUFFER\_MANAGEMENT\_AREA data structure that is indicated by the IA32\_DS\_AREA register. The layout of the DS\_BUFFER\_MANAGEMENT\_AREA for Goldmont Plus is shown in Figure 18-40. When a counter generates a PEBS records, the appropriate counter reset values will be loaded into that counter. In the above example where general purpose counter 2 and fixed function counter 0 generated the PEBS record, general purpose counter 2 would be reloaded with the value contained in PEBS GP Counter 2 Reset (offset 50H) and fixed function counter 0 would be reloaded with the value contained in PEBS Fixed Counter 0 Reset (offset 80H).




Figure 18-40. PEBS Programming Environment

### 18.5.4.2 Reduced Skid PEBS

Goldmont Plus microarchitecture processors supports the Reduced Skid PEBS feature described in Section 18.5.3.1.2 on the IA32\_PMC0 counter. Although Goldmont Plus adds support for generating PEBS records for precise events on the other general-purpose and fixed-function performance counters, those counters do not support the Reduced Skid PEBS feature.

## 18.6 PERFORMANCE MONITORING (LEGACY INTEL PROCESSORS)

### 18.6.1 Performance Monitoring (Intel® Core™ Solo and Intel® Core™ Duo Processors)

In Intel Core Solo and Intel Core Duo processors, non-architectural performance monitoring events are programmed using the same facilities (see Figure 18-1) used for architectural performance events.

Non-architectural performance events use event select values that are model-specific. Event mask (Umask) values are also specific to event logic units. Some microarchitectural conditions detectable by a Umask value may have specificity related to processor topology (see Section 8.6, “Detecting Hardware Multi-Threading Support and Topology,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*). As a result, the unit mask field (for example, IA32\_PERFEVTSELx[bits 15:8]) may contain sub-fields that specify topology information of processor cores.

The sub-field layout within the Umask field may support two-bit encoding that qualifies the relationship between a microarchitectural condition and the originating core. This data is shown in Table 18-61. The two-bit encoding for core-specificity is only supported for a subset of Umask values (see Chapter 19, “Performance Monitoring Events”) and for Intel Core Duo processors. Such events are referred to as core-specific events.

**Table 18-61. Core Specificity Encoding within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |             |
|-----------------------|-------------|
| Bit 15:14 Encoding    | Description |
| 11B                   | All cores   |
| 10B                   | Reserved    |
| 01B                   | This core   |
| 00B                   | Reserved    |

Some microarchitectural conditions allow detection specificity only at the boundary of physical processors. Some bus events belong to this category, providing specificity between the originating physical processor (a bus agent) versus other agents on the bus. Sub-field encoding for agent specificity is shown in Table 18-62.

**Table 18-62. Agent Specificity Encoding within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |                    |
|-----------------------|--------------------|
| Bit 13 Encoding       | Description        |
| 0                     | This agent         |
| 1                     | Include all agents |

Some microarchitectural conditions are detectable only from the originating core. In such cases, unit mask does not support core-specificity or agent-specificity encodings. These are referred to as core-only conditions.

Some microarchitectural conditions allow detection specificity that includes or excludes the action of hardware prefetches. A two-bit encoding may be supported to qualify hardware prefetch actions. Typically, this applies only to some L2 or bus events. The sub-field encoding for hardware prefetch qualification is shown in Table 18-63.

**Table 18-63. HW Prefetch Qualification Encoding within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |                           |
|-----------------------|---------------------------|
| Bit 13:12 Encoding    | Description               |
| 11B                   | All inclusive             |
| 10B                   | Reserved                  |
| 01B                   | Hardware prefetch only    |
| 00B                   | Exclude hardware prefetch |

Some performance events may (a) support none of the three event-specific qualification encodings (b) may support core-specificity and agent specificity simultaneously (c) or may support core-specificity and hardware prefetch qualification simultaneously. Agent-specificity and hardware prefetch qualification are mutually exclusive.

In addition, some L2 events permit qualifications that distinguish cache coherent states. The sub-field definition for cache coherency state qualification is shown in Table 18-64. If no bits in the MESI qualification sub-field are set for an event that requires setting MESI qualification bits, the event count will not increment.

**Table 18-64. MESI Qualification Definitions within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |                        |
|-----------------------|------------------------|
| Bit Position 11:8     | Description            |
| Bit 11                | Counts modified state  |
| Bit 10                | Counts exclusive state |
| Bit 9                 | Counts shared state    |
| Bit 8                 | Counts Invalid state   |

## 18.6.2 Performance Monitoring (Processors Based on Intel® Core™ Microarchitecture)

In addition to architectural performance monitoring, processors based on the Intel Core microarchitecture support non-architectural performance monitoring events.

Architectural performance events can be collected using general-purpose performance counters. Non-architectural performance events can be collected using general-purpose performance counters (coupled with two IA32\_PERFEVTSELx MSRs for detailed event configurations), or fixed-function performance counters (see Section 18.6.2.1). IA32\_PERFEVTSELx MSRs are architectural; their layout is shown in Figure 18-1. Starting with Intel Core 2 processor T 7700, fixed-function performance counters and associated counter control and status MSR becomes part of architectural performance monitoring version 2 facilities (see also Section 18.2.2).

Non-architectural performance events in processors based on Intel Core microarchitecture use event select values that are model-specific. Valid event mask (Umask) bits are listed in Chapter 19. The UMASK field may contain sub-fields identical to those listed in Table 18-61, Table 18-62, Table 18-63, and Table 18-64. One or more of these sub-fields may apply to specific events on an event-by-event basis. Details are listed in Table 19-25 in Chapter 19, "Performance Monitoring Events."

In addition, the UMASK field may also contain a sub-field that allows detection specificity related to snoop responses. Bits of the snoop response qualification sub-field are defined in Table 18-65.

**Table 18-65. Bus Snoop Qualification Definitions within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |               |
|-----------------------|---------------|
| Bit Position 11:8     | Description   |
| Bit 11                | HITM response |
| Bit 10                | Reserved      |

**Table 18-65. Bus Snoop Qualification Definitions within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |                |
|-----------------------|----------------|
| Bit Position 11:8     | Description    |
| Bit 9                 | HIT response   |
| Bit 8                 | CLEAN response |

There are also non-architectural events that support qualification of different types of snoop operation. The corresponding bit field for snoop type qualification are listed in Table 18-66.

**Table 18-66. Snoop Type Qualification Definitions within a Non-Architectural Umask**

| IA32_PERFEVTSELx MSRs |              |
|-----------------------|--------------|
| Bit Position 9:8      | Description  |
| Bit 9                 | CMP2I snoops |
| Bit 8                 | CMP2S snoops |

No more than one sub-field of MESI, snoop response, and snoop type qualification sub-fields can be supported in a performance event.

#### NOTE

Software must write known values to the performance counters prior to enabling the counters. The content of general-purpose counters and fixed-function counters are undefined after INIT or RESET.

### 18.6.2.1 Fixed-function Performance Counters

Processors based on Intel Core microarchitecture provide three fixed-function performance counters. Bits beyond the width of the fixed counter are reserved and must be written as zeros. Model-specific fixed-function performance counters on processors that support Architectural Perfmon version 1 are 40 bits wide.

Each of the fixed-function counter is dedicated to count a pre-defined performance monitoring events. See Table 18-2 for details of the PMC addresses and what these events count.

Programming the fixed-function performance counters does not involve any of the IA32\_PERFEVTSELx MSRs, and does not require specifying any event masks. Instead, the MSR MSR\_PERF\_FIXED\_CTR\_CTRL provides multiple sets of 4-bit fields; each 4-bit field controls the operation of a fixed-function performance counter (PMC). See Figures 18-41. Two sub-fields are defined for each control. See Figure 18-41; bit fields are:

- Enable field (low 2 bits in each 4-bit control)** — When bit 0 is set, performance counting is enabled in the corresponding fixed-function performance counter to increment when the target condition associated with the architecture performance event occurs at ring 0.  
 When bit 1 is set, performance counting is enabled in the corresponding fixed-function performance counter to increment when the target condition associated with the architecture performance event occurs at ring greater than 0.  
 Writing 0 to both bits stops the performance counter. Writing 11B causes the counter to increment irrespective of privilege levels.

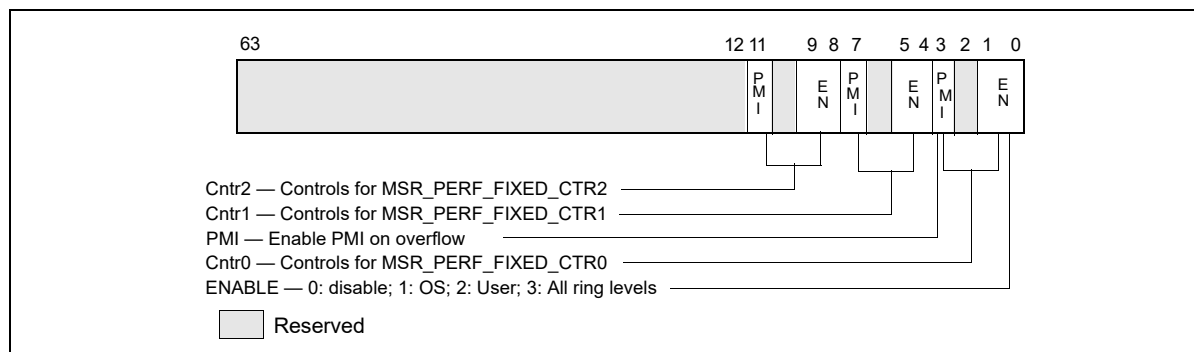


Figure 18-41. Layout of MSR\_PERF\_FIXED\_CTRL MSR

- **PMI field (fourth bit in each 4-bit control)** — When set, the logical processor generates an exception through its local APIC on overflow condition of the respective fixed-function counter.

### 18.6.2.2 Global Counter Control Facilities

Processors based on Intel Core microarchitecture provides simplified performance counter control that simplifies the most frequent operations in programming performance events, i.e. enabling/disabling event counting and checking the status of counter overflows. This is done by the following three MSR:

- MSR\_PERF\_GLOBAL\_CTRL enables/disables event counting for all or any combination of fixed-function PMCs (MSR\_PERF\_FIXED\_CTRLx) or general-purpose PMCs via a single WRMSR.
- MSR\_PERF\_GLOBAL\_STATUS allows software to query counter overflow conditions on any combination of fixed-function PMCs (MSR\_PERF\_FIXED\_CTRLx) or general-purpose PMCs via a single RDMSR.
- MSR\_PERF\_GLOBAL\_OVF\_CTRL allows software to clear counter overflow conditions on any combination of fixed-function PMCs (MSR\_PERF\_FIXED\_CTRLx) or general-purpose PMCs via a single WRMSR.

MSR\_PERF\_GLOBAL\_CTRL MSR provides single-bit controls to enable counting in each performance counter (see Figure 18-42). Each enable bit in MSR\_PERF\_GLOBAL\_CTRL is AND'ed with the enable bits for all privilege levels in the respective IA32\_PERFEVTSELx or MSR\_PERF\_FIXED\_CTRL\_CTRL MSRs to start/stop the counting of respective counters. Counting is enabled if the AND'ed results is true; counting is disabled when the result is false.




Figure 18-42. Layout of MSR\_PERF\_GLOBAL\_CTRL MSR

MSR\_PERF\_GLOBAL\_STATUS MSR provides single-bit status used by software to query the overflow condition of each performance counter. MSR\_PERF\_GLOBAL\_STATUS[bit 62] indicates overflow conditions of the DS area data buffer. MSR\_PERF\_GLOBAL\_STATUS[bit 63] provides a CondChgd bit to indicate changes to the state of performance monitoring hardware (see Figure 18-43). A value of 1 in bits 34:32, 1, 0 indicates an overflow condition has

occurred in the associated counter.




Figure 18-43. Layout of MSR\_PERF\_GLOBAL\_STATUS MSR

When a performance counter is configured for PEBS, an overflow condition in the counter will arm PEBS. On the subsequent event following overflow, the processor will generate a PEBS event. On a PEBS event, the processor will perform bounds checks based on the parameters defined in the DS Save Area (see Section 17.4.9). Upon successful bounds checks, the processor will store the data record in the defined buffer area, clear the counter overflow status, and reload the counter. If the bounds checks fail, the PEBS will be skipped entirely. In the event that the PEBS buffer fills up, the processor will set the OvfBuffer bit in MSR\_PERF\_GLOBAL\_STATUS.

MSR\_PERF\_GLOBAL\_OVF\_CTL MSR allows software to clear overflow the indicators for general-purpose or fixed-function counters via a single WRMSR (see Figure 18-44). Clear overflow indications when:

- Setting up new values in the event select and/or UMASK field for counting or interrupt-based event sampling.
- Reloading counter values to continue collecting next sample.
- Disabling event counting or interrupt-based event sampling.




Figure 18-44. Layout of MSR\_PERF\_GLOBAL\_OVF\_CTL MSR

### 18.6.2.3 At-Retirement Events

Many non-architectural performance events are impacted by the speculative nature of out-of-order execution. A subset of non-architectural performance events on processors based on Intel Core microarchitecture are enhanced with a tagging mechanism (similar to that found in Intel NetBurst<sup>®</sup> microarchitecture) that exclude contributions that arise from speculative execution. The at-retirement events available in processors based on Intel Core microarchitecture does not require special MSR programming control (see Section 18.6.3.6, "At-Retirement Counting"), but is limited to IA32\_PMC0. See Table 18-67 for a list of events available to processors based on Intel Core microarchitecture.



**Table 18-67. At-Retirement Performance Events for Intel Core Microarchitecture**

| Event Name                     | UMask | Event Select |
|--------------------------------|-------|--------------|
| ITLB_MISS_RETIRED              | 00H   | C9H          |
| MEM_LOAD_RETIRED.L1D_MISS      | 01H   | CBH          |
| MEM_LOAD_RETIRED.L1D_LINE_MISS | 02H   | CBH          |
| MEM_LOAD_RETIRED.L2_MISS       | 04H   | CBH          |
| MEM_LOAD_RETIRED.L2_LINE_MISS  | 08H   | CBH          |
| MEM_LOAD_RETIRED.DTLB_MISS     | 10H   | CBH          |

#### 18.6.2.4 Processor Event Based Sampling (PEBS)

Processors based on Intel Core microarchitecture also support processor event based sampling (PEBS). This feature was introduced by processors based on Intel NetBurst microarchitecture.

PEBS uses a debug store mechanism and a performance monitoring interrupt to store a set of architectural state information for the processor. The information provides architectural state of the instruction executed after the instruction that caused the event (See Section 18.6.2.4.2 and Section 17.4.9).

In cases where the same instruction causes BTS and PEBS to be activated, PEBS is processed before BTS are processed. The PMI request is held until the processor completes processing of PEBS and BTS.

For processors based on Intel Core microarchitecture, precise events that can be used with PEBS are listed in Table 18-68. The procedure for detecting availability of PEBS is the same as described in Section 18.6.3.8.1.

**Table 18-68. PEBS Performance Events for Intel Core Microarchitecture**

| Event Name                     | UMask | Event Select |
|--------------------------------|-------|--------------|
| INSTR_RETIRED.ANY_P            | 00H   | C0H          |
| X87_OPS_RETIRED.ANY            | FEH   | C1H          |
| BR_INST_RETIRED.MISPRED        | 00H   | C5H          |
| SIMD_INST_RETIRED.ANY          | 1FH   | C7H          |
| MEM_LOAD_RETIRED.L1D_MISS      | 01H   | CBH          |
| MEM_LOAD_RETIRED.L1D_LINE_MISS | 02H   | CBH          |
| MEM_LOAD_RETIRED.L2_MISS       | 04H   | CBH          |
| MEM_LOAD_RETIRED.L2_LINE_MISS  | 08H   | CBH          |
| MEM_LOAD_RETIRED.DTLB_MISS     | 10H   | CBH          |

##### 18.6.2.4.1 Setting up the PEBS Buffer

For processors based on Intel Core microarchitecture, PEBS is available using IA32\_PMC0 only. Use the following procedure to set up the processor and IA32\_PMC0 counter for PEBS:

1. Set up the precise event buffering facilities. Place values in the precise event buffer base, precise event index, precise event absolute maximum, precise event interrupt threshold, and precise event counter reset fields of the DS buffer management area. In processors based on Intel Core microarchitecture, PEBS records consist of 64-bit address entries. See Figure 17-8 to set up the precise event records buffer in memory.
2. Enable PEBS. Set the Enable PEBS on PMC0 flag (bit 0) in IA32\_PEBS\_ENABLE MSR.
3. Set up the IA32\_PMC0 performance counter and IA32\_PERFEVTSEL0 for an event listed in Table 18-68.

##### 18.6.2.4.2 PEBS Record Format

The PEBS record format may be extended across different processor implementations. The IA32\_PERF\_CAPABILITIES MSR defines a mechanism for software to handle the evolution of PEBS record format in processors that support architectural performance monitoring with version id equals 2 or higher. The bit fields of

IA32\_PERF\_CAPABILITIES are defined in Table 2-2 of Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*. The relevant bit fields that governs PEBS are:

- PEBSTrap [bit 6]: When set, PEBS recording is trap-like. After the PEBS-enabled counter has overflowed, PEBS record is recorded for the next PEBS-able event at the completion of the sampled instruction causing the PEBS event. When clear, PEBS recording is fault-like. The PEBS record is recorded before the sampled instruction causing the PEBS event.
- PEBSaveArchRegs [bit 7]: When set, PEBS will save architectural register and state information according to the encoded value of the PEBSRecordFormat field. When clear, only the return instruction pointer and flags are recorded. On processors based on Intel Core microarchitecture, this bit is always 1
- PEBSRecordFormat [bits 11:8]: Valid encodings are:
  - 0000B: Only general-purpose registers, instruction pointer and RFLAGS registers are saved in each PEBS record (seeSection 18.6.3.8).
  - 0001B: PEBS record includes additional information of IA32\_PERF\_GLOBAL\_STATUS and load latency data. (seeSection 18.3.1.1.1).
  - 0010B: PEBS record includes additional information of IA32\_PERF\_GLOBAL\_STATUS, load latency data, and TSX tuning information. (seeSection 18.3.6.2).
  - 0011B: PEBS record includes additional information of load latency data, TSX tuning information, TSC data, and the applicable counter field replaces IA32\_PERF\_GLOBAL\_STATUS at offset 90H. (see Section 18.3.8.1.1).

### 18.6.2.4.3 Writing a PEBS Interrupt Service Routine

The PEBS facilities share the same interrupt vector and interrupt service routine (called the DS ISR) with the Interrupt-based event sampling and BTS facilities. To handle PEBS interrupts, PEBS handler code must be included in the DS ISR. See Section 17.4.9.1, “64 Bit Format of the DS Save Area,” for guidelines when writing the DS ISR.

The service routine can query MSR\_PERF\_GLOBAL\_STATUS to determine which counter(s) caused of overflow condition. The service routine should clear overflow indicator by writing to MSR\_PERF\_GLOBAL\_OVF\_CTL.

A comparison of the sequence of requirements to program PEBS for processors based on Intel Core and Intel NetBurst microarchitectures is listed in Table 18-69.

**Table 18-69. Requirements to Program PEBS**

|                                      | For Processors based on Intel Core microarchitecture  | For Processors based on Intel NetBurst microarchitecture |
|--------------------------------------|---|--|
| Verify PEBS support of processor/OS. | <ul style="list-style-type: none"> <li>▪ IA32_MISC_ENABLE.EMON_AVAILABE (bit 7) is set.</li> <li>▪ IA32_MISC_ENABLE.PEBS_UNAVAILABE (bit 12) is clear.</li> </ul>   |  |
| Ensure counters are in disabled.     | On initial set up or changing event configurations, write MSR_PERF_GLOBAL_CTRL MSR (38FH) with 0.<br>On subsequent entries: <ul style="list-style-type: none"> <li>▪ Clear all counters if “Counter Freeze on PMI” is not enabled.</li> <li>▪ If IA32_DebugCTL.Freeze is enabled, counters are automatically disabled.</li> </ul> Counters MUST be stopped before writing. <sup>1</sup> | Optional   |
| Disable PEBS.                        | Clear ENABLE PMCO bit in IA32_PEBS_ENABLE MSR (3F1H).   | Optional   |
| Check overflow conditions.           | Check MSR_PERF_GLOBAL_STATUS MSR (38EH) handle any overflow conditions.   | Check OVF flag of each CCCR for overflow condition       |
| Clear overflow status.               | Clear MSR_PERF_GLOBAL_STATUS MSR (38EH) using IA32_PERF_GLOBAL_OVF_CTRL MSR (390H).   | Clear OVF flag of each CCCR.                             |
| Write “sample-after” values.         | Configure the counter(s) with the sample after value.   |  |

**Table 18-69. Requirements to Program PEBS (Contd.)**

|   | For Processors based on Intel Core microarchitecture  | For Processors based on Intel NetBurst microarchitecture   |
|---|---|--|
| Configure specific counter configuration MSR. | <ul style="list-style-type: none"> <li>▪ Set local enable bit 22 - 1.</li> <li>▪ Do NOT set local counter PMI/INT bit, bit 20 - 0.</li> <li>▪ Event programmed must be PEBS capable.</li> </ul> | <ul style="list-style-type: none"> <li>▪ Set appropriate OVF_PMI bits - 1.</li> <li>▪ Only CCCR for MSR_IQ_COUNTER4 support PEBS.</li> </ul> |
| Allocate buffer for PEBS states.              | Allocate a buffer in memory for the precise information.  |  |
| Program the IA32_DS_AREA MSR.                 | Program the IA32_DS_AREA MSR.   |  |
| Configure the PEBS buffer management records. | Configure the PEBS buffer management records in the DS buffer management area.  |  |
| Configure/Enable PEBS.                        | Set Enable PMCO bit in IA32_PEBS_ENABLE MSR (3F1H).   | Configure MSR_PEBS_ENABLE, MSR_PEBS_MATRIX_VERT and MSR_PEBS_MATRIX_HORZ as needed.  |
| Enable counters.                              | Set Enable bits in MSR_PERF_GLOBAL_CTRL MSR (38FH).   | Set each CCCR enable bit 12 - 1.   |

**NOTES:**

1. Counters read while enabled are not guaranteed to be precise with event counts that occur in timing proximity to the RDMSR.

**18.6.2.4.4 Re-configuring PEBS Facilities**

When software needs to reconfigure PEBS facilities, it should allow a quiescent period between stopping the prior event counting and setting up a new PEBS event. The quiescent period is to allow any latent residual PEBS records to complete its capture at their previously specified buffer address (provided by IA32\_DS\_AREA).

**18.6.3 Performance Monitoring (Processors Based on Intel NetBurst® Microarchitecture)**

The performance monitoring mechanism provided in processors based on Intel NetBurst microarchitecture is different from that provided in the P6 family and Pentium processors. While the general concept of selecting, filtering, counting, and reading performance events through the WRMSR, RDMSR, and RDPMC instructions is unchanged, the setup mechanism and MSR layouts are incompatible with the P6 family and Pentium processor mechanisms. Also, the RDPMC instruction has been extended to support faster reading of counters and to read all performance counters available in processors based on Intel NetBurst microarchitecture.

The event monitoring mechanism consists of the following facilities:

- The IA32\_MISC\_ENABLE MSR, which indicates the availability in an Intel 64 or IA-32 processor of the performance monitoring and processor event-based sampling (PEBS) facilities.
- Event selection control (ESCR) MSRs for selecting events to be monitored with specific performance counters. The number available differs by family and model (43 to 45).
- 18 performance counter MSRs for counting events.
- 18 counter configuration control (CCCR) MSRs, with one CCCR associated with each performance counter. CCCRs sets up an associated performance counter for a specific method of counting.
- A debug store (DS) save area in memory for storing PEBS records.
- The IA32\_DS\_AREA MSR, which establishes the location of the DS save area.
- The debug store (DS) feature flag (bit 21) returned by the CPUID instruction, which indicates the availability of the DS mechanism.
- The MSR\_PEBS\_ENABLE MSR, which enables the PEBS facilities and replay tagging used in at-retirement event counting.
- A set of predefined events and event metrics that simplify the setting up of the performance counters to count specific events.

Table 18-70 lists the performance counters and their associated CCCRs, along with the ESCRs that select events to be counted for each performance counter. Predefined event metrics and events are listed in Chapter 19, "Perfor-

mance Monitoring Events.”

**Table 18-70. Performance Counter MSRs and Associated CCCR and ESCR MSRs (Processors Based on Intel NetBurst Microarchitecture)**

| Counter            |     |      | CCCR            |      | ESCR            |     |      |
|--------------------|-----|------|-----------------|------|-----------------|-----|------|
| Name               | No. | Addr | Name            | Addr | Name            | No. | Addr |
| MSR_BPU_COUNTER0   | 0   | 300H | MSR_BPU_CCCR0   | 360H | MSR_BSU_ESCR0   | 7   | 3A0H |
|                    |     |      |                 |      | MSR_FSB_ESCR0   | 6   | 3A2H |
|                    |     |      |                 |      | MSR_MOB_ESCR0   | 2   | 3AAH |
|                    |     |      |                 |      | MSR_PMH_ESCR0   | 4   | 3ACH |
|                    |     |      |                 |      | MSR_BPU_ESCR0   | 0   | 3B2H |
|                    |     |      |                 |      | MSR_IS_ESCR0    | 1   | 3B4H |
|                    |     |      |                 |      | MSR_ITLB_ESCR0  | 3   | 3B6H |
|                    |     |      |                 |      | MSR_IX_ESCR0    | 5   | 3C8H |
| MSR_BPU_COUNTER1   | 1   | 301H | MSR_BPU_CCCR1   | 361H | MSR_BSU_ESCR0   | 7   | 3A0H |
|                    |     |      |                 |      | MSR_FSB_ESCR0   | 6   | 3A2H |
|                    |     |      |                 |      | MSR_MOB_ESCR0   | 2   | 3AAH |
|                    |     |      |                 |      | MSR_PMH_ESCR0   | 4   | 3ACH |
|                    |     |      |                 |      | MSR_BPU_ESCR0   | 0   | 3B2H |
|                    |     |      |                 |      | MSR_IS_ESCR0    | 1   | 3B4H |
|                    |     |      |                 |      | MSR_ITLB_ESCR0  | 3   | 3B6H |
|                    |     |      |                 |      | MSR_IX_ESCR0    | 5   | 3C8H |
| MSR_BPU_COUNTER2   | 2   | 302H | MSR_BPU_CCCR2   | 362H | MSR_BSU_ESCR1   | 7   | 3A1H |
|                    |     |      |                 |      | MSR_FSB_ESCR1   | 6   | 3A3H |
|                    |     |      |                 |      | MSR_MOB_ESCR1   | 2   | 3ABH |
|                    |     |      |                 |      | MSR_PMH_ESCR1   | 4   | 3ADH |
|                    |     |      |                 |      | MSR_BPU_ESCR1   | 0   | 3B3H |
|                    |     |      |                 |      | MSR_IS_ESCR1    | 1   | 3B5H |
|                    |     |      |                 |      | MSR_ITLB_ESCR1  | 3   | 3B7H |
|                    |     |      |                 |      | MSR_IX_ESCR1    | 5   | 3C9H |
| MSR_BPU_COUNTER3   | 3   | 303H | MSR_BPU_CCCR3   | 363H | MSR_BSU_ESCR1   | 7   | 3A1H |
|                    |     |      |                 |      | MSR_FSB_ESCR1   | 6   | 3A3H |
|                    |     |      |                 |      | MSR_MOB_ESCR1   | 2   | 3ABH |
|                    |     |      |                 |      | MSR_PMH_ESCR1   | 4   | 3ADH |
|                    |     |      |                 |      | MSR_BPU_ESCR1   | 0   | 3B3H |
|                    |     |      |                 |      | MSR_IS_ESCR1    | 1   | 3B5H |
|                    |     |      |                 |      | MSR_ITLB_ESCR1  | 3   | 3B7H |
|                    |     |      |                 |      | MSR_IX_ESCR1    | 5   | 3C9H |
| MSR_MS_COUNTER0    | 4   | 304H | MSR_MS_CCCR0    | 364H | MSR_MS_ESCR0    | 0   | 3C0H |
|                    |     |      |                 |      | MSR_TBPU_ESCR0  | 2   | 3C2H |
|                    |     |      |                 |      | MSR_TC_ESCR0    | 1   | 3C4H |
| MSR_MS_COUNTER1    | 5   | 305H | MSR_MS_CCCR1    | 365H | MSR_MS_ESCR0    | 0   | 3C0H |
|                    |     |      |                 |      | MSR_TBPU_ESCR0  | 2   | 3C2H |
|                    |     |      |                 |      | MSR_TC_ESCR0    | 1   | 3C4H |
| MSR_MS_COUNTER2    | 6   | 306H | MSR_MS_CCCR2    | 366H | MSR_MS_ESCR1    | 0   | 3C1H |
|                    |     |      |                 |      | MSR_TBPU_ESCR1  | 2   | 3C3H |
|                    |     |      |                 |      | MSR_TC_ESCR1    | 1   | 3C5H |
| MSR_MS_COUNTER3    | 7   | 307H | MSR_MS_CCCR3    | 367H | MSR_MS_ESCR1    | 0   | 3C1H |
|                    |     |      |                 |      | MSR_TBPU_ESCR1  | 2   | 3C3H |
|                    |     |      |                 |      | MSR_TC_ESCR1    | 1   | 3C5H |
| MSR_FLAME_COUNTER0 | 8   | 308H | MSR_FLAME_CCCR0 | 368H | MSR_FIRM_ESCR0  | 1   | 3A4H |
|                    |     |      |                 |      | MSR_FLAME_ESCR0 | 0   | 3A6H |
|                    |     |      |                 |      | MSR_DAC_ESCR0   | 5   | 3A8H |
|                    |     |      |                 |      | MSR_SAA_T_ESCR0 | 2   | 3AEH |
|                    |     |      |                 |      | MSR_U2L_ESCR0   | 3   | 3B0H |

**Table 18-70. Performance Counter MSRs and Associated CCCR and ESCR MSRs (Processors Based on Intel NetBurst Microarchitecture) (Contd.)**

| Counter            |     |      | CCCR            |      | ESCR  |                                 |  |
|--------------------|-----|------|-----------------|------|---|---------------------------------|--|
| Name               | No. | Addr | Name            | Addr | Name  | No.                             | Addr   |
| MSR_FLAME_COUNTER1 | 9   | 309H | MSR_FLAME_CCCR1 | 369H | MSR_FIRM_ESCR0<br>MSR_FLAME_ESCR0<br>MSR_DAC_ESCR0<br>MSR_SAA_T_ESCR0<br>MSR_U2L_ESCR0  | 1<br>0<br>5<br>2<br>3           | 3A4H<br>3A6H<br>3A8H<br>3AEH<br>3B0H                 |
| MSR_FLAME_COUNTER2 | 10  | 30AH | MSR_FLAME_CCCR2 | 36AH | MSR_FIRM_ESCR1<br>MSR_FLAME_ESCR1<br>MSR_DAC_ESCR1<br>MSR_SAA_T_ESCR1<br>MSR_U2L_ESCR1  | 1<br>0<br>5<br>2<br>3           | 3A5H<br>3A7H<br>3A9H<br>3AFH<br>3B1H                 |
| MSR_FLAME_COUNTER3 | 11  | 30BH | MSR_FLAME_CCCR3 | 36BH | MSR_FIRM_ESCR1<br>MSR_FLAME_ESCR1<br>MSR_DAC_ESCR1<br>MSR_SAA_T_ESCR1<br>MSR_U2L_ESCR1  | 1<br>0<br>5<br>2<br>3           | 3A5H<br>3A7H<br>3A9H<br>3AFH<br>3B1H                 |
| MSR_IQ_COUNTER0    | 12  | 30CH | MSR_IQ_CCCR0    | 36CH | MSR_CRU_ESCR0<br>MSR_CRU_ESCR2<br>MSR_CRU_ESCR4<br>MSR_IQ_ESCR0 <sup>1</sup><br>MSR_RAT_ESCR0<br>MSR_SSU_ESCR0<br>MSR_ALF_ESCR0 | 4<br>5<br>6<br>0<br>2<br>3<br>1 | 3B8H<br>3CCH<br>3E0H<br>3BAH<br>3BCH<br>3BEH<br>3CAH |
| MSR_IQ_COUNTER1    | 13  | 30DH | MSR_IQ_CCCR1    | 36DH | MSR_CRU_ESCR0<br>MSR_CRU_ESCR2<br>MSR_CRU_ESCR4<br>MSR_IQ_ESCR0 <sup>1</sup><br>MSR_RAT_ESCR0<br>MSR_SSU_ESCR0<br>MSR_ALF_ESCR0 | 4<br>5<br>6<br>0<br>2<br>3<br>1 | 3B8H<br>3CCH<br>3E0H<br>3BAH<br>3BCH<br>3BEH<br>3CAH |
| MSR_IQ_COUNTER2    | 14  | 30EH | MSR_IQ_CCCR2    | 36EH | MSR_CRU_ESCR1<br>MSR_CRU_ESCR3<br>MSR_CRU_ESCR5<br>MSR_IQ_ESCR1 <sup>1</sup><br>MSR_RAT_ESCR1<br>MSR_ALF_ESCR1                  | 4<br>5<br>6<br>0<br>2<br>1      | 3B9H<br>3CDH<br>3E1H<br>3BBH<br>3BDH<br>3CBH         |
| MSR_IQ_COUNTER3    | 15  | 30FH | MSR_IQ_CCCR3    | 36FH | MSR_CRU_ESCR1<br>MSR_CRU_ESCR3<br>MSR_CRU_ESCR5<br>MSR_IQ_ESCR1 <sup>1</sup><br>MSR_RAT_ESCR1<br>MSR_ALF_ESCR1                  | 4<br>5<br>6<br>0<br>2<br>1      | 3B9H<br>3CDH<br>3E1H<br>3BBH<br>3BDH<br>3CBH         |
| MSR_IQ_COUNTER4    | 16  | 310H | MSR_IQ_CCCR4    | 370H | MSR_CRU_ESCR0<br>MSR_CRU_ESCR2<br>MSR_CRU_ESCR4<br>MSR_IQ_ESCR0 <sup>1</sup><br>MSR_RAT_ESCR0<br>MSR_SSU_ESCR0<br>MSR_ALF_ESCR0 | 4<br>5<br>6<br>0<br>2<br>3<br>1 | 3B8H<br>3CCH<br>3E0H<br>3BAH<br>3BCH<br>3BEH<br>3CAH |
| MSR_IQ_COUNTER5    | 17  | 311H | MSR_IQ_CCCR5    | 371H | MSR_CRU_ESCR1<br>MSR_CRU_ESCR3<br>MSR_CRU_ESCR5<br>MSR_IQ_ESCR1 <sup>1</sup><br>MSR_RAT_ESCR1<br>MSR_ALF_ESCR1                  | 4<br>5<br>6<br>0<br>2<br>1      | 3B9H<br>3CDH<br>3E1H<br>3BBH<br>3BDH<br>3CBH         |

**NOTES:**

1. MSR\_IQ\_ESCR0 and MSR\_IQ\_ESCR1 are available only on early processor builds (family 0FH, models 01H-02H). These MSRs are not available on later versions.

The types of events that can be counted with these performance monitoring facilities are divided into two classes: non-retirement events and at-retirement events.

- Non-retirement events (see Table 19-31) are events that occur any time during instruction execution (such as bus transactions or cache transactions).
- At-retirement events (see Table 19-32) are events that are counted at the retirement stage of instruction execution, which allows finer granularity in counting events and capturing machine state.

The at-retirement counting mechanism includes facilities for tagging  $\mu$ ops that have encountered a particular performance event during instruction execution. Tagging allows events to be sorted between those that occurred on an execution path that resulted in architectural state being committed at retirement as well as events that occurred on an execution path where the results were eventually cancelled and never committed to architectural state (such as, the execution of a mispredicted branch).

The Pentium 4 and Intel Xeon processor performance monitoring facilities support the three usage models described below. The first two models can be used to count both non-retirement and at-retirement events; the third model is used to count a subset of at-retirement events:

- **Event counting** — A performance counter is configured to count one or more types of events. While the counter is counting, software reads the counter at selected intervals to determine the number of events that have been counted between the intervals.
- **Interrupt-based event sampling** — A performance counter is configured to count one or more types of events and to generate an interrupt when it overflows. To trigger an overflow, the counter is preset to a modulus value that will cause the counter to overflow after a specific number of events have been counted.

When the counter overflows, the processor generates a performance monitoring interrupt (PMI). The interrupt service routine for the PMI then records the return instruction pointer (RIP), resets the modulus, and restarts the counter. Code performance can be analyzed by examining the distribution of RIPs with a tool like the VTune™ Performance Analyzer.

- **Processor event-based sampling (PEBS)** — In PEBS, the processor writes a record of the architectural state of the processor to a memory buffer after the counter overflows. The records of architectural state provide additional information for use in performance tuning. Processor-based event sampling can be used to count only a subset of at-retirement events. PEBS captures more precise processor state information compared to interrupt based event sampling, because the latter need to use the interrupt service routine to re-construct the architectural states of processor.


The following sections describe the MSRs and data structures used for performance monitoring in the Pentium 4 and Intel Xeon processors.

### 18.6.3.1 ESCR MSRs

The 45 ESCR MSRs (see Table 18-70) allow software to select specific events to be countered. Each ESCR is usually associated with a pair of performance counters (see Table 18-70) and each performance counter has several ESCRs associated with it (allowing the events counted to be selected from a variety of events).

Figure 18-45 shows the layout of an ESCR MSR. The functions of the flags and fields are:

- **USR flag, bit 2** — When set, events are counted when the processor is operating at a current privilege level (CPL) of 1, 2, or 3. These privilege levels are generally used by application code and unprotected operating system code.
- **OS flag, bit 3** — When set, events are counted when the processor is operating at CPL of 0. This privilege level is generally reserved for protected operating system code. (When both the OS and USR flags are set, events are counted at all privilege levels.)



**Figure 18-45. Event Selection Control Register (ESCR) for Pentium 4 and Intel Xeon Processors without Intel HT Technology Support**

- **Tag enable, bit 4** — When set, enables tagging of  $\mu$ ops to assist in at-retirement event counting; when clear, disables tagging. See Section 18.6.3.6, “At-Retirement Counting.”
- **Tag value field, bits 5 through 8** — Selects a tag value to associate with a  $\mu$ op to assist in at-retirement event counting.
- **Event mask field, bits 9 through 24** — Selects events to be counted from the event class selected with the event select field.
- **Event select field, bits 25 through 30** — Selects a class of events to be counted. The events within this class that are counted are selected with the event mask field.

When setting up an ESCR, the event select field is used to select a specific class of events to count, such as retired branches. The event mask field is then used to select one or more of the specific events within the class to be counted. For example, when counting retired branches, four different events can be counted: branch not taken predicted, branch not taken mispredicted, branch taken predicted, and branch taken mispredicted. The OS and USR flags allow counts to be enabled for events that occur when operating system code and/or application code are being executed. If neither the OS nor USR flag is set, no events will be counted.

The ESCRs are initialized to all 0s on reset. The flags and fields of an ESCR are configured by writing to the ESCR using the WRMSR instruction. Table 18-70 gives the addresses of the ESCR MSRs.

Writing to an ESCR MSR does not enable counting with its associated performance counter; it only selects the event or events to be counted. The CCCR for the selected performance counter must also be configured. Configuration of the CCCR includes selecting the ESCR and enabling the counter.

### 18.6.3.2 Performance Counters

The performance counters in conjunction with the counter configuration control registers (CCCRs) are used for filtering and counting the events selected by the ESCRs. Processors based on Intel NetBurst microarchitecture provide 18 performance counters organized into 9 pairs. A pair of performance counters is associated with a particular subset of events and ESCR's (see Table 18-70). The counter pairs are partitioned into four groups:

- The BPU group, includes two performance counter pairs:
  - MSR\_BPU\_COUNTER0 and MSR\_BPU\_COUNTER1.
  - MSR\_BPU\_COUNTER2 and MSR\_BPU\_COUNTER3.
- The MS group, includes two performance counter pairs:
  - MSR\_MS\_COUNTER0 and MSR\_MS\_COUNTER1.
  - MSR\_MS\_COUNTER2 and MSR\_MS\_COUNTER3.
- The FLAME group, includes two performance counter pairs:
  - MSR\_FLAME\_COUNTER0 and MSR\_FLAME\_COUNTER1.

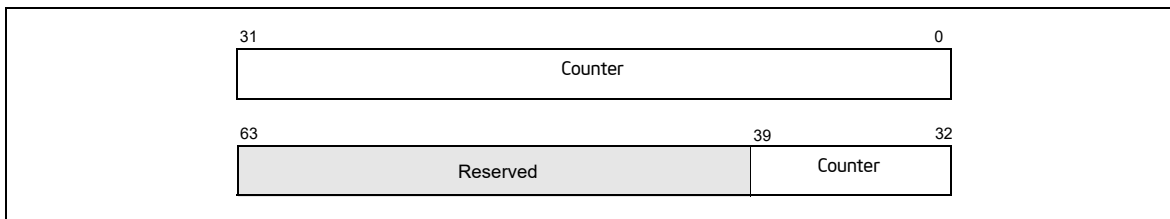
- MSR\_FLAME\_COUNTER2 and MSR\_FLAME\_COUNTER3.
- The IQ group, includes three performance counter pairs:
  - MSR\_IQ\_COUNTER0 and MSR\_IQ\_COUNTER1.
  - MSR\_IQ\_COUNTER2 and MSR\_IQ\_COUNTER3.
  - MSR\_IQ\_COUNTER4 and MSR\_IQ\_COUNTER5.

The MSR\_IQ\_COUNTER4 counter in the IQ group provides support for the PEBS.

Alternate counters in each group can be cascaded: the first counter in one pair can start the first counter in the second pair and vice versa. A similar cascading is possible for the second counters in each pair. For example, within the BPU group of counters, MSR\_BPU\_COUNTER0 can start MSR\_BPU\_COUNTER2 and vice versa, and MSR\_BPU\_COUNTER1 can start MSR\_BPU\_COUNTER3 and vice versa (see Section 18.6.3.5.6, “Cascading Counters”). The cascade flag in the CCCR register for the performance counter enables the cascading of counters.

Each performance counter is 40-bits wide (see Figure 18-46). The RDPMC instruction is intended to allow reading of either the full counter-width (40-bits) or the low 32-bits of the counter. Reading the low 32-bits is faster than reading the full counter width and is appropriate in situations where the count is small enough to be contained in 32 bits.

The RDPMC instruction can be used by programs or procedures running at any privilege level and in virtual-8086 mode to read these counters. The PCE flag in control register CR4 (bit 8) allows the use of this instruction to be restricted to only programs and procedures running at privilege level 0.



**Figure 18-46. Performance Counter (Pentium 4 and Intel Xeon Processors)**

The RDPMC instruction is not serializing or ordered with other instructions. Thus, it does not necessarily wait until all previous instructions have been executed before reading the counter. Similarly, subsequent instructions may begin execution before the RDPMC instruction operation is performed.

Only the operating system, executing at privilege level 0, can directly manipulate the performance counters, using the RDMSR and WRMSR instructions. A secure operating system would clear the PCE flag during system initialization to disable direct user access to the performance-monitoring counters, but provide a user-accessible programming interface that emulates the RDPMC instruction.

Some uses of the performance counters require the counters to be preset before counting begins (that is, before the counter is enabled). This can be accomplished by writing to the counter using the WRMSR instruction. To set a counter to a specified number of counts before overflow, enter a 2s complement negative integer in the counter. The counter will then count from the preset value up to -1 and overflow. Writing to a performance counter in a Pentium 4 or Intel Xeon processor with the WRMSR instruction causes all 40 bits of the counter to be written.


### 18.6.3.3 CCCR MSRs

Each of the 18 performance counters has one CCCR MSR associated with it (see Table 18-70). The CCCRs control the filtering and counting of events as well as interrupt generation. Figure 18-47 shows the layout of an CCCR MSR. The functions of the flags and fields are as follows:

- **Enable flag, bit 12** — When set, enables counting; when clear, the counter is disabled. This flag is cleared on reset.
- **ESCR select field, bits 13 through 15** — Identifies the ESCR to be used to select events to be counted with the counter associated with the CCCR.



- **Compare flag, bit 18** — When set, enables filtering of the event count; when clear, disables filtering. The filtering method is selected with the threshold, complement, and edge flags.
- **Complement flag, bit 19** — Selects how the incoming event count is compared with the threshold value. When set, event counts that are less than or equal to the threshold value result in a single count being delivered to the performance counter; when clear, counts greater than the threshold value result in a count being delivered to the performance counter (see Section 18.6.3.5.2, “Filtering Events”). The complement flag is not active unless the compare flag is set.
- **Threshold field, bits 20 through 23** — Selects the threshold value to be used for comparisons. The processor examines this field only when the compare flag is set, and uses the complement flag setting to determine the type of threshold comparison to be made. The useful range of values that can be entered in this field depend on the type of event being counted (see Section 18.6.3.5.2, “Filtering Events”).
- **Edge flag, bit 24** — When set, enables rising edge (false-to-true) edge detection of the threshold comparison output for filtering event counts; when clear, rising edge detection is disabled. This flag is active only when the compare flag is set.



**Figure 18-47. Counter Configuration Control Register (CCCR)**

- **FORCE\_OVF flag, bit 25** — When set, forces a counter overflow on every counter increment; when clear, overflow only occurs when the counter actually overflows.
- **OVF\_PMI flag, bit 26** — When set, causes a performance monitor interrupt (PMI) to be generated when the counter overflows occurs; when clear, disables PMI generation. Note that the PMI is generated on the next event count after the counter has overflowed.
- **Cascade flag, bit 30** — When set, enables counting on one counter of a counter pair when its alternate counter in the other the counter pair in the same counter group overflows (see Section 18.6.3.2, “Performance Counters,” for further details); when clear, disables cascading of counters.
- **OVF flag, bit 31** — Indicates that the counter has overflowed when set. This flag is a sticky flag that must be explicitly cleared by software.

The CCCRs are initialized to all 0s on reset.

The events that an enabled performance counter actually counts are selected and filtered by the following flags and fields in the ESCR and CCCR registers and in the qualification order given:

1. The event select and event mask fields in the ESCR select a class of events to be counted and one or more event types within the class, respectively.

2. The OS and USR flags in the ESCR selected the privilege levels at which events will be counted.
3. The ESCR select field of the CCCR selects the ESCR. Since each counter has several ESCRs associated with it, one ESCR must be chosen to select the classes of events that may be counted.
4. The compare and complement flags and the threshold field of the CCCR select an optional threshold to be used in qualifying an event count.
5. The edge flag in the CCCR allows events to be counted only on rising-edge transitions.

The qualification order in the above list implies that the filtered output of one "stage" forms the input for the next. For instance, events filtered using the privilege level flags can be further qualified by the compare and complement flags and the threshold field, and an event that matched the threshold criteria, can be further qualified by edge detection.

The uses of the flags and fields in the CCCRs are discussed in greater detail in Section 18.6.3.5, "Programming the Performance Counters for Non-Retirement Events."

### 18.6.3.4 Debug Store (DS) Mechanism

The debug store (DS) mechanism was introduced with processors based on Intel NetBurst microarchitecture to allow various types of information to be collected in memory-resident buffers for use in debugging and tuning programs. The DS mechanism can be used to collect two types of information: branch records and processor event-based sampling (PEBS) records. The availability of the DS mechanism in a processor is indicated with the DS feature flag (bit 21) returned by the CPUID instruction.

See Section 17.4.5, "Branch Trace Store (BTS)," and Section 18.6.3.8, "Processor Event-Based Sampling (PEBS)," for a description of these facilities. Records collected with the DS mechanism are saved in the DS save area. See Section 17.4.9, "BTS and DS Save Area."

### 18.6.3.5 Programming the Performance Counters for Non-Retirement Events

The basic steps to program a performance counter and to count events include the following:

1. Select the event or events to be counted.
2. For each event, select an ESCR that supports the event using the values in the ESCR restrictions row in Table 19-31, Chapter 19.
3. Match the CCCR Select value and ESCR name in Table 19-31 to a value listed in Table 18-70; select a CCCR and performance counter.
4. Set up an ESCR for the specific event or events to be counted and the privilege levels at which they are to be counted.
5. Set up the CCCR for the performance counter by selecting the ESCR and the desired event filters.
6. Set up the CCCR for optional cascading of event counts, so that when the selected counter overflows its alternate counter starts.
7. Set up the CCCR to generate an optional performance monitor interrupt (PMI) when the counter overflows. If PMI generation is enabled, the local APIC must be set up to deliver the interrupt to the processor and a handler for the interrupt must be in place.
8. Enable the counter to begin counting.

### 18.6.3.5.1 Selecting Events to Count

Table 19-32 in Chapter 19 lists a set of at-retirement events for processors based on Intel NetBurst microarchitecture. For each event listed in Table 19-32, setup information is provided. Table 18-71 gives an example of one of the events.

**Table 18-71. Event Example**

| Event Name                           | Event Parameters         | Parameter Value                              | Description  |
|--------------------------------------|--------------------------|--|--|
| branch_retired                       |                          |  | Counts the retirement of a branch. Specify one or more mask bits to select any combination of branch taken, not-taken, predicted and mispredicted. |
|                                      | ESCR restrictions        | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3               | See Table 15-3 for the addresses of the ESCR MSRs.   |
|                                      | Counter numbers per ESCR | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17       | The counter numbers associated with each ESCR are provided. The performance counters and corresponding CCCRs can be obtained from Table 15-3.      |
|                                      | ESCR Event Select        | 06H  | ESCR[31:25]  |
|                                      | ESCR Event Mask          | Bit 0: MMNP<br>1: MMNM<br>2: MMTP<br>3: MMTM | ESCR[24:9]<br>Branch Not-taken Predicted<br>Branch Not-taken Mispredicted<br>Branch Taken Predicted<br>Branch Taken Mispredicted                   |
|                                      | CCCR Select              | 05H  | CCCR[15:13]  |
|                                      | Event Specific Notes     |  | P6: EMON_BR_INST_RETIRED   |
|                                      | Can Support PEBS         | No   |  |
| Requires Additional MSRs for Tagging | No                       |  |  |

For Table 19-31 and Table 19-32, Chapter 19, the name of the event is listed in the Event Name column and parameters that define the event and other information are listed in the Event Parameters column. The Parameter Value and Description columns give specific parameters for the event and additional description information. Entries in the Event Parameters column are described below.

- **ESCR restrictions** — Lists the ESCRs that can be used to program the event. Typically only one ESCR is needed to count an event.
- **Counter numbers per ESCR** — Lists which performance counters are associated with each ESCR. Table 18-70 gives the name of the counter and CCCR for each counter number. Typically only one counter is needed to count the event.
- **ESCR event select** — Gives the value to be placed in the event select field of the ESCR to select the event.
- **ESCR event mask** — Gives the value to be placed in the Event Mask field of the ESCR to select sub-events to be counted. The parameter value column defines the documented bits with relative bit position offset starting from 0, where the absolute bit position of relative offset 0 is bit 9 of the ESCR. All undocumented bits are reserved and should be set to 0.
- **CCCR select** — Gives the value to be placed in the ESCR select field of the CCCR associated with the counter to select the ESCR to be used to define the event. This value is not the address of the ESCR; it is the number of the ESCR from the Number column in Table 18-70.
- **Event specific notes** — Gives additional information about the event, such as the name of the same or a similar event defined for the P6 family processors.
- **Can support PEBS** — Indicates if PEBS is supported for the event (only supplied for at-retirement events listed in Table 19-32.)
- **Requires additional MSR for tagging** — Indicates which if any additional MSRs must be programmed to count the events (only supplied for the at-retirement events listed in Table 19-32.)

**NOTE**

The performance-monitoring events listed in Chapter 19, "Performance Monitoring Events," are intended to be used as guides for performance tuning. The counter values reported are not guaranteed to be absolutely accurate and should be used as a relative guide for tuning. Known discrepancies are documented where applicable.

The following procedure shows how to set up a performance counter for basic counting; that is, the counter is set up to count a specified event indefinitely, wrapping around whenever it reaches its maximum count. This procedure is continued through the following four sections.

Using information in Table 19-31, Chapter 19, an event to be counted can be selected as follows:

1. Select the event to be counted.
2. Select the ESCR to be used to select events to be counted from the ESCRs field.
3. Select the number of the counter to be used to count the event from the Counter Numbers Per ESCR field.
4. Determine the name of the counter and the CCCR associated with the counter, and determine the MSR addresses of the counter, CCCR, and ESCR from Table 18-70.
5. Use the WRMSR instruction to write the ESCR Event Select and ESCR Event Mask values into the appropriate fields in the ESCR. At the same time set or clear the USR and OS flags in the ESCR as desired.
6. Use the WRMSR instruction to write the CCCR Select value into the appropriate field in the CCCR.

**NOTE**

Typically all the fields and flags of the CCCR will be written with one WRMSR instruction; however, in this procedure, several WRMSR writes are used to more clearly demonstrate the uses of the various CCCR fields and flags.

This setup procedure is continued in the next section, Section 18.6.3.5.2, "Filtering Events."

**18.6.3.5.2 Filtering Events**

Each counter receives up to 4 input lines from the processor hardware from which it is counting events. The counter treats these inputs as binary inputs (input 0 has a value of 1, input 1 has a value of 2, input 2 has a value of 4, and input 3 has a value of 8). When a counter is enabled, it adds this binary input value to the counter value on each clock cycle. For each clock cycle, the value added to the counter can then range from 0 (no event) to 15.

For many events, only the 0 input line is active, so the counter is merely counting the clock cycles during which the 0 input is asserted. However, for some events two or more input lines are used. Here, the counter's threshold setting can be used to filter events. The compare, complement, threshold, and edge fields control the filtering of counter increments by input value.


If the compare flag is set, then a "greater than" or a "less than or equal to" comparison of the input value vs. a threshold value can be made. The complement flag selects "less than or equal to" (flag set) or "greater than" (flag clear). The threshold field selects a threshold value of from 0 to 15. For example, if the complement flag is cleared and the threshold field is set to 6, then any input value of 7 or greater on the 4 inputs to the counter will cause the counter to be incremented by 1, and any value less than 7 will cause an increment of 0 (or no increment) of the counter. Conversely, if the complement flag is set, any value from 0 to 6 will increment the counter and any value from 7 to 15 will not increment the counter. Note that when a threshold condition has been satisfied, the input to the counter is always 1, not the input value that is presented to the threshold filter.

The edge flag provides further filtering of the counter inputs when a threshold comparison is being made. The edge flag is only active when the compare flag is set. When the edge flag is set, the resulting output from the threshold filter (a value of 0 or 1) is used as an input to the edge filter. Each clock cycle, the edge filter examines the last and current input values and sends a count to the counter only when it detects a "rising edge" event; that is, a false-to-true transition. Figure 18-48 illustrates rising edge filtering.

The following procedure shows how to configure a CCCR to filter events using the threshold filter and the edge filter. This procedure is a continuation of the setup procedure introduced in Section 18.6.3.5.1, “Selecting Events to Count.”

7. (Optional) To set up the counter for threshold filtering, use the WRMSR instruction to write values in the CCCR compare and complement flags and the threshold field:
  - Set the compare flag.
  - Set or clear the complement flag for less than or equal to or greater than comparisons, respectively.
  - Enter a value from 0 to 15 in the threshold field.
8. (Optional) Select rising edge filtering by setting the CCCR edge flag.

This setup procedure is continued in the next section, Section 18.6.3.5.3, “Starting Event Counting.”



**Figure 18-48. Effects of Edge Filtering**

### 18.6.3.5.3 Starting Event Counting

Event counting by a performance counter can be initiated in either of two ways. The typical way is to set the enable flag in the counter’s CCCR. Following the instruction to set the enable flag, event counting begins and continues until it is stopped (see Section 18.6.3.5.5, “Halting Event Counting”).

The following procedural step shows how to start event counting. This step is a continuation of the setup procedure introduced in Section 18.6.3.5.2, “Filtering Events.”

9. To start event counting, use the WRMSR instruction to set the CCCR enable flag for the performance counter.

This setup procedure is continued in the next section, Section 18.6.3.5.4, “Reading a Performance Counter’s Count.”

The second way that a counter can be started by using the cascade feature. Here, the overflow of one counter automatically starts its alternate counter (see Section 18.6.3.5.6, “Cascading Counters”).

### 18.6.3.5.4 Reading a Performance Counter’s Count

Performance counters can be read using either the RDPMC or RDMSR instructions. The enhanced functions of the RDPMC instruction (including fast read) are described in Section 18.6.3.2, “Performance Counters.” These instructions can be used to read a performance counter while it is counting or when it is stopped.

The following procedural step shows how to read the event counter. This step is a continuation of the setup procedure introduced in Section 18.6.3.5.3, “Starting Event Counting.”

10. To read a performance counter’s current event count, execute the RDPMC instruction with the counter number obtained from Table 18-70 used as an operand.

This setup procedure is continued in the next section, Section 18.6.3.5.5, “Halting Event Counting.”

### 18.6.3.5.5 Halting Event Counting

After a performance counter has been started (enabled), it continues counting indefinitely. If the counter overflows (goes one count past its maximum count), it wraps around and continues counting. When the counter wraps

around, it sets its OVF flag to indicate that the counter has overflowed. The OVF flag is a sticky flag that indicates that the counter has overflowed at least once since the OVF bit was last cleared.

To halt counting, the CCCR enable flag for the counter must be cleared.

The following procedural step shows how to stop event counting. This step is a continuation of the setup procedure introduced in Section 18.6.3.5.4, “Reading a Performance Counter’s Count.”

11. To stop event counting, execute a WRMSR instruction to clear the CCCR enable flag for the performance counter.

To halt a cascaded counter (a counter that was started when its alternate counter overflowed), either clear the Cascade flag in the cascaded counter’s CCCR MSR or clear the OVF flag in the alternate counter’s CCCR MSR.

### 18.6.3.5.6 Cascading Counters

As described in Section 18.6.3.2, “Performance Counters,” eighteen performance counters are implemented in pairs. Nine pairs of counters and associated CCCRs are further organized as four blocks: BPU, MS, FLAME, and IQ (see Table 18-70). The first three blocks contain two pairs each. The IQ block contains three pairs of counters (12 through 17) with associated CCCRs (MSR\_IQ\_CCCR0 through MSR\_IQ\_CCCR5).

The first 8 counter pairs (0 through 15) can be programmed using ESCRs to detect performance monitoring events. Pairs of ESCRs in each of the four blocks allow many different types of events to be counted. The cascade flag in the CCCR MSR allows nested monitoring of events to be performed by cascading one counter to a second counter located in another pair in the same block (see Figure 18-47 for the location of the flag).

Counters 0 and 1 form the first pair in the BPU block. Either counter 0 or 1 can be programmed to detect an event via MSR\_MOB\_ESCR0. Counters 0 and 2 can be cascaded in any order, as can counters 1 and 3. It’s possible to set up 4 counters in the same block to cascade on two pairs of independent events. The pairing described also applies to subsequent blocks. Since the IQ PUB has two extra counters, cascading operates somewhat differently if 16 and 17 are involved. In the IQ block, counter 16 can only be cascaded from counter 14 (not from 12); counter 14 cannot be cascaded from counter 16 using the CCCR cascade bit mechanism. Similar restrictions apply to counter 17.

#### Example 18-1. Counting Events

Assume a scenario where counter X is set up to count 200 occurrences of event A; then counter Y is set up to count 400 occurrences of event B. Each counter is set up to count a specific event and overflow to the next counter. In the above example, counter X is preset for a count of -200 and counter Y for a count of -400; this setup causes the counters to overflow on the 200th and 400th counts respectively.

Continuing this scenario, counter X is set up to count indefinitely and wraparound on overflow. This is described in the basic performance counter setup procedure that begins in Section 18.6.3.5.1, “Selecting Events to Count.” Counter Y is set up with the cascade flag in its associated CCCR MSR set to 1 and its enable flag set to 0.

To begin the nested counting, the enable bit for the counter X is set. Once enabled, counter X counts until it overflows. At this point, counter Y is automatically enabled and begins counting. Thus counter X overflows after 200 occurrences of event A. Counter Y then starts, counting 400 occurrences of event B before overflowing. When performance counters are cascaded, the counter Y would typically be set up to generate an interrupt on overflow. This is described in Section 18.6.3.5.8, “Generating an Interrupt on Overflow.”

The cascading counters mechanism can be used to count a single event. The counting begins on one counter then continues on the second counter after the first counter overflows. This technique doubles the number of event counts that can be recorded, since the contents of the two counters can be added together.

### 18.6.3.5.7 EXTENDED CASCADING

Extended cascading is a model-specific feature in the Intel NetBurst microarchitecture with CPUID DisplayFamily\_DisplayModel 0F\_02, 0F\_03, 0F\_04, 0F\_06. This feature uses bit 11 in CCCRs associated with the IQ

block. See Table 18-72.

**Table 18-72. CCR Names and Bit Positions**

| CCCR Name:Bit Position | Bit Name     | Description                               |
|------------------------|--------------|---|
| MSR_IQ_CCCR1 2:11      | Reserved     |   |
| MSR_IQ_CCCR0:11        | CASCNT4INT00 | Allow counter 4 to cascade into counter 0 |
| MSR_IQ_CCCR3:11        | CASCNT5INT03 | Allow counter 5 to cascade into counter 3 |
| MSR_IQ_CCCR4:11        | CASCNT5INT04 | Allow counter 5 to cascade into counter 4 |
| MSR_IQ_CCCR5:11        | CASCNT4INT05 | Allow counter 4 to cascade into counter 5 |

The extended cascading feature can be adapted to the Interrupt based sampling usage model for performance monitoring. However, it is known that performance counters do not generate PMI in cascade mode or extended cascade mode due to an erratum. This erratum applies to processors with CPUID DisplayFamily\_DisplayModel signature of 0F\_02. For processors with CPUID DisplayFamily\_DisplayModel signature of 0F\_00 and 0F\_01, the erratum applies to processors with stepping encoding greater than 09H.

Counters 16 and 17 in the IQ block are frequently used in processor event-based sampling or at-retirement counting of events indicating a stalled condition in the pipeline. Neither counter 16 or 17 can initiate the cascading of counter pairs using the cascade bit in a CCCR.

Extended cascading permits performance monitoring tools to use counters 16 and 17 to initiate cascading of two counters in the IQ block. Extended cascading from counter 16 and 17 is conceptually similar to cascading other counters, but instead of using CASCADE bit of a CCCR, one of the four CASCNTxINT0y bits is used.

#### Example 18-2. Scenario for Extended Cascading

A usage scenario for extended cascading is to sample instructions retired on logical processor 1 after the first 4096 instructions retired on logical processor 0. A procedure to program extended cascading in this scenario is outlined below:

1. Write the value 0 to counter 12.
2. Write the value 04000603H to MSR\_CRU\_ESCR0 (corresponding to selecting the NBOGNTAG and NBOGTAG event masks with qualification restricted to logical processor 1).
3. Write the value 04038800H to MSR\_IQ\_CCCR0. This enables CASCNT4INT00 and OVF\_PMI. An ISR can sample on instruction addresses in this case (do not set ENABLE, or CASCADE).
4. Write the value FFFF000H into counter 16.1.
5. Write the value 0400060CH to MSR\_CRU\_ESCR2 (corresponding to selecting the NBOGNTAG and NBOGTAG event masks with qualification restricted to logical processor 0).
6. Write the value 00039000H to MSR\_IQ\_CCCR4 (set ENABLE bit, but not OVF\_PMI).

Another use for cascading is to locate stalled execution in a multithreaded application. Assume MOB replays in thread B cause thread A to stall. Getting a sample of the stalled execution in this scenario could be accomplished by:

1. Set up counter B to count MOB replays on thread B.
2. Set up counter A to count resource stalls on thread A; set its force overflow bit and the appropriate CASCNTx-INT0y bit.
3. Use the performance monitoring interrupt to capture the program execution data of the stalled thread.

#### 18.6.3.5.8 Generating an Interrupt on Overflow

Any performance counter can be configured to generate a performance monitor interrupt (PMI) if the counter overflows. The PMI interrupt service routine can then collect information about the state of the processor or program

when overflow occurred. This information can then be used with a tool like the Intel® VTune™ Performance Analyzer to analyze and tune program performance.

To enable an interrupt on counter overflow, the OVR\_PMI flag in the counter's associated CCCR MSR must be set. When overflow occurs, a PMI is generated through the local APIC. (Here, the performance counter entry in the local vector table [LVT] is set up to deliver the interrupt generated by the PMI to the processor.)

The PMI service routine can use the OVF flag to determine which counter overflowed when multiple counters have been configured to generate PMIs. Also, note that these processors mask PMIs upon receiving an interrupt. Clear this condition before leaving the interrupt handler.

When generating interrupts on overflow, the performance counter being used should be preset to value that will cause an overflow after a specified number of events are counted plus 1. The simplest way to select the preset value is to write a negative number into the counter, as described in Section 18.6.3.5.6, "Cascading Counters." Here, however, if an interrupt is to be generated after 100 event counts, the counter should be preset to minus 100 plus 1 (-100 + 1), or -99. The counter will then overflow after it counts 99 events and generate an interrupt on the next (100th) event counted. The difference of 1 for this count enables the interrupt to be generated immediately after the selected event count has been reached, instead of waiting for the overflow to be propagation through the counter.

Because of latency in the microarchitecture between the generation of events and the generation of interrupts on overflow, it is sometimes difficult to generate an interrupt close to an event that caused it. In these situations, the FORCE\_OVF flag in the CCCR can be used to improve reporting. Setting this flag causes the counter to overflow on every counter increment, which in turn triggers an interrupt after every counter increment.

#### 18.6.3.5.9 Counter Usage Guideline

There are some instances where the user must take care to configure counting logic properly, so that it is not powered down. To use any ESCR, even when it is being used just for tagging, (any) one of the counters that the particular ESCR (or its paired ESCR) can be connected to should be enabled. If this is not done, 0 counts may result. Likewise, to use any counter, there must be some event selected in a corresponding ESCR (other than no\_event, which generally has a select value of 0).

#### 18.6.3.6 At-Retirement Counting

At-retirement counting provides a means counting only events that represent work committed to architectural state and ignoring work that was performed speculatively and later discarded.

One example of this speculative activity is branch prediction. When a branch misprediction occurs, the results of instructions that were decoded and executed down the mispredicted path are canceled. If a performance counter was set up to count all executed instructions, the count would include instructions whose results were canceled as well as those whose results committed to architectural state.

To provide finer granularity in event counting in these situations, the performance monitoring facilities provided in the Pentium 4 and Intel Xeon processors provide a mechanism for tagging events and then counting only those tagged events that represent committed results. This mechanism is called "at-retirement counting."

Tables 19-32 through 19-36 list predefined at-retirement events and event metrics that can be used to for tagging events when using at retirement counting. The following terminology is used in describing at-retirement counting:

- **Bogus, non-bogus, retire** — In at-retirement event descriptions, the term "bogus" refers to instructions or  $\mu$ ops that must be canceled because they are on a path taken from a mispredicted branch. The terms "retired" and "non-bogus" refer to instructions or  $\mu$ ops along the path that results in committed architectural state changes as required by the program being executed. Thus instructions and  $\mu$ ops are either bogus or non-bogus, but not both. Several of the Pentium 4 and Intel Xeon processors' performance monitoring events (such as, Instruction\_Retired and Uops\_Retired in Table 19-32) can count instructions or  $\mu$ ops that are retired based on the characterization of bogus" versus non-bogus.
- **Tagging** — Tagging is a means of marking  $\mu$ ops that have encountered a particular performance event so they can be counted at retirement. During the course of execution, the same event can happen more than once per  $\mu$ op and a direct count of the event would not provide an indication of how many  $\mu$ ops encountered that event. The tagging mechanisms allow a  $\mu$ op to be tagged once during its lifetime and thus counted once at retirement. The retired suffix is used for performance metrics that increment a count once per  $\mu$ op, rather than once per



event. For example, a  $\mu$ op may encounter a cache miss more than once during its life time, but a “Miss Retired” metric (that counts the number of retired  $\mu$ ops that encountered a cache miss) will increment only once for that  $\mu$ op. A “Miss Retired” metric would be useful for characterizing the performance of the cache hierarchy for a particular instruction sequence. Details of various performance metrics and how these can be constructed using the Pentium 4 and Intel Xeon processors performance events are provided in the *Intel Pentium 4 Processor Optimization Reference Manual* (see Section 1.4, “Related Literature”).

- **Replay** — To maximize performance for the common case, the Intel NetBurst microarchitecture aggressively schedules  $\mu$ ops for execution before all the conditions for correct execution are guaranteed to be satisfied. In the event that all of these conditions are not satisfied,  $\mu$ ops must be reissued. The mechanism that the Pentium 4 and Intel Xeon processors use for this reissuing of  $\mu$ ops is called replay. Some examples of replay causes are cache misses, dependence violations, and unforeseen resource constraints. In normal operation, some number of replays is common and unavoidable. An excessive number of replays is an indication of a performance problem.
- **Assist** — When the hardware needs the assistance of microcode to deal with some event, the machine takes an assist. One example of this is an underflow condition in the input operands of a floating-point operation. The hardware must internally modify the format of the operands in order to perform the computation. Assists clear the entire machine of  $\mu$ ops before they begin and are costly.

#### 18.6.3.6.1 Using At-Retirement Counting

Processors based on Intel NetBurst microarchitecture allow counting both events and  $\mu$ ops that encountered a specified event. For a subset of the at-retirement events listed in Table 19-32, a  $\mu$ op may be tagged when it encounters that event. The tagging mechanisms can be used in Interrupt-based event sampling, and a subset of these mechanisms can be used in PEBS. There are four independent tagging mechanisms, and each mechanism uses a different event to count  $\mu$ ops tagged with that mechanism:

- **Front-end tagging** — This mechanism pertains to the tagging of  $\mu$ ops that encountered front-end events (for example, trace cache and instruction counts) and are counted with the `Front_end_event` event.
- **Execution tagging** — This mechanism pertains to the tagging of  $\mu$ ops that encountered execution events (for example, instruction types) and are counted with the `Execution_Event` event.
- **Replay tagging** — This mechanism pertains to tagging of  $\mu$ ops whose retirement is replayed (for example, a cache miss) and are counted with the `Replay_event` event. Branch mispredictions are also tagged with this mechanism.
- **No tags** — This mechanism does not use tags. It uses the `Instr_retired` and the `Uops_retired` events.

Each tagging mechanism is independent from all others; that is, a  $\mu$ op that has been tagged using one mechanism will not be detected with another mechanism’s tagged- $\mu$ op detector. For example, if  $\mu$ ops are tagged using the front-end tagging mechanisms, the `Replay_event` will not count those as tagged  $\mu$ ops unless they are also tagged using the replay tagging mechanism. However, execution tags allow up to four different types of  $\mu$ ops to be counted at retirement through execution tagging.

The independence of tagging mechanisms does not hold when using PEBS. When using PEBS, only one tagging mechanism should be used at a time.

Certain kinds of  $\mu$ ops that cannot be tagged, including I/O, uncacheable and locked accesses, returns, and far transfers.

Table 19-32 lists the performance monitoring events that support at-retirement counting: specifically the `Front_end_event`, `Execution_event`, `Replay_event`, `Inst_retired` and `Uops_retired` events. The following sections describe the tagging mechanisms for using these events to tag  $\mu$ op and count tagged  $\mu$ ops.

#### 18.6.3.6.2 Tagging Mechanism for `Front_end_event`

The `Front_end_event` counts  $\mu$ ops that have been tagged as encountering any of the following events:

- **$\mu$ op decode events** — Tagging  $\mu$ ops for  $\mu$ op decode events requires specifying bits in the `ESCR` associated with the performance-monitoring event, `Uop_type`.
- **Trace cache events** — Tagging  $\mu$ ops for trace cache events may require specifying certain bits in the `MSR_TC_PRECISE_EVENT` MSR (see Table 19-34).

Table 19-32 describes the `Front_end_event` and Table 19-34 describes metrics that are used to set up a `Front_end_event` count.

The MSR's specified in the Table 19-32 that are supported by the front-end tagging mechanism must be set and one or both of the `NBOGUS` and `BOGUS` bits in the `Front_end_event` event mask must be set to count events. None of the events currently supported requires the use of the `MSR_TC_PRECISE_EVENT` MSR.

### 18.6.3.6.3 Tagging Mechanism For Execution\_event

Table 19-32 describes the `Execution_event` and Table 19-35 describes metrics that are used to set up an `Execution_event` count.

The execution tagging mechanism differs from other tagging mechanisms in how it causes tagging. One *upstream* ESCR is used to specify an event to detect and to specify a tag value (bits 5 through 8) to identify that event. A second *downstream* ESCR is used to detect  $\mu$ ops that have been tagged with that tag value identifier using `Execution_event` for the event selection.

The upstream ESCR that counts the event must have its tag enable flag (bit 4) set and must have an appropriate tag value mask entered in its tag value field. The 4-bit tag value mask specifies which of tag bits should be set for a particular  $\mu$ op. The value selected for the tag value should coincide with the event mask selected in the downstream ESCR. For example, if a tag value of 1 is set, then the event mask of `NBOGUS0` should be enabled, correspondingly in the downstream ESCR. The downstream ESCR detects and counts tagged  $\mu$ ops. The normal (not tag value) mask bits in the downstream ESCR specify which tag bits to count. If any one of the tag bits selected by the mask is set, the related counter is incremented by one. This mechanism is summarized in the Table 19-35 metrics that are supported by the execution tagging mechanism. The tag enable and tag value bits are irrelevant for the downstream ESCR used to select the `Execution_event`.

The four separate tag bits allow the user to simultaneously but distinctly count up to four execution events at retirement. (This applies for interrupt-based event sampling. There are additional restrictions for PEBS as noted in Section 18.6.3.8.3, "Setting Up the PEBS Buffer.") It is also possible to detect or count combinations of events by setting multiple tag value bits in the upstream ESCR or multiple mask bits in the downstream ESCR. For example, use a tag value of 3H in the upstream ESCR and use `NBOGUS0/NBOGUS1` in the downstream ESCR event mask.

### 18.6.3.7 Tagging Mechanism for Replay\_event

Table 19-32 describes the `Replay_event` and Table 19-36 describes metrics that are used to set up an `Replay_event` count.

The replay mechanism enables tagging of  $\mu$ ops for a subset of all replays before retirement. Use of the replay mechanism requires selecting the type of  $\mu$ op that may experience the replay in the `MSR_PEBS_MATRIX_VERT` MSR and selecting the type of event in the `MSR_PEBS_ENABLE` MSR. Replay tagging must also be enabled with the `UOP_Tag` flag (bit 24) in the `MSR_PEBS_ENABLE` MSR.

The Table 19-36 lists the metrics that are support the replay tagging mechanism and the at-retirement events that use the replay tagging mechanism, and specifies how the appropriate MSR's need to be configured. The replay tags defined in Table A-5 also enable Processor Event-Based Sampling (PEBS, see Section 17.4.9). Each of these replay tags can also be used in normal sampling by not setting Bit 24 nor Bit 25 in `IA_32_PEBS_ENABLE_MSR`. Each of these metrics requires that the `Replay_Event` (see Table 19-32) be used to count the tagged  $\mu$ ops.

### 18.6.3.8 Processor Event-Based Sampling (PEBS)

The debug store (DS) mechanism in processors based on Intel NetBurst microarchitecture allow two types of information to be collected for use in debugging and tuning programs: PEBS records and BTS records. See Section 17.4.5, "Branch Trace Store (BTS)," for a description of the BTS mechanism.

PEBS permits the saving of precise architectural information associated with one or more performance events in the precise event records buffer, which is part of the DS save area (see Section 17.4.9, "BTS and DS Save Area"). To use this mechanism, a counter is configured to overflow after it has counted a preset number of events. After the counter overflows, the processor copies the current state of the general-purpose and `EFLAGS` registers and instruction pointer into a record in the precise event records buffer. The processor then resets the count in the performance counter and restarts the counter. When the precise event records buffer is nearly full, an interrupt is

generated, allowing the precise event records to be saved. A circular buffer is not supported for precise event records.

PEBS is supported only for a subset of the at-retirement events: Execution\_event, Front\_end\_event, and Replay\_event. Also, PEBS can only be carried out using the one performance counter, the MSR\_IQ\_COUNTER4 MSR.

In processors based on Intel Core microarchitecture, a similar PEBS mechanism is also supported using IA32\_PMC0 and IA32\_PERFEVTSEL0 MSRs (See Section 18.6.2.4).

#### 18.6.3.8.1 Detection of the Availability of the PEBS Facilities

The DS feature flag (bit 21) returned by the CPUID instruction indicates (when set) the availability of the DS mechanism in the processor, which supports the PEBS (and BTS) facilities. When this bit is set, the following PEBS facilities are available:

- The PEBS\_UNAVAILABLE flag in the IA32\_MISC\_ENABLE MSR indicates (when clear) the availability of the PEBS facilities, including the MSR\_PEBS\_ENABLE MSR.
- The enable PEBS flag (bit 24) in the MSR\_PEBS\_ENABLE MSR allows PEBS to be enabled (set) or disabled (clear).
- The IA32\_DS\_AREA MSR can be programmed to point to the DS save area.

#### 18.6.3.8.2 Setting Up the DS Save Area

Section 17.4.9.2, "Setting Up the DS Save Area," describes how to set up and enable the DS save area. This procedure is common for PEBS and BTS.

#### 18.6.3.8.3 Setting Up the PEBS Buffer

Only the MSR\_IQ\_COUNTER4 performance counter can be used for PEBS. Use the following procedure to set up the processor and this counter for PEBS:

1. Set up the precise event buffering facilities. Place values in the precise event buffer base, precise event index, precise event absolute maximum, and precise event interrupt threshold, and precise event counter reset fields of the DS buffer management area (see Figure 17-5) to set up the precise event records buffer in memory.
2. Enable PEBS. Set the Enable PEBS flag (bit 24) in MSR\_PEBS\_ENABLE MSR.
3. Set up the MSR\_IQ\_COUNTER4 performance counter and its associated CCCR and one or more ESCRs for PEBS as described in Tables 19-32 through 19-36.

#### 18.6.3.8.4 Writing a PEBS Interrupt Service Routine

The PEBS facilities share the same interrupt vector and interrupt service routine (called the DS ISR) with the non-precise event-based sampling and BTS facilities. To handle PEBS interrupts, PEBS handler code must be included in the DS ISR. See Section 17.4.9.5, "Writing the DS Interrupt Service Routine," for guidelines for writing the DS ISR.

#### 18.6.3.8.5 Other DS Mechanism Implications

The DS mechanism is not available in the SMM. It is disabled on transition to the SMM mode. Similarly the DS mechanism is disabled on the generation of a machine check exception and is cleared on processor RESET and INIT.

The DS mechanism is available in real address mode.

#### 18.6.3.9 Operating System Implications

The DS mechanism can be used by the operating system as a debugging extension to facilitate failure analysis. When using this facility, a 25 to 30 times slowdown can be expected due to the effects of the trace store occurring on every taken branch.

Depending upon intended usage, the instruction pointers that are part of the branch records or the PEBS records need to have an association with the corresponding process. One solution requires the ability for the DS specific operating system module to be chained to the context switch. A separate buffer can then be maintained for each process of interest and the MSR pointing to the configuration area saved and setup appropriately on each context switch.

If the BTS facility has been enabled, then it must be disabled and state stored on transition of the system to a sleep state in which processor context is lost. The state must be restored on return from the sleep state.

It is required that an interrupt gate be used for the DS interrupt as opposed to a trap gate to prevent the generation of an endless interrupt loop.

Pages that contain buffers must have mappings to the same physical address for all processes/logical processors, such that any change to CR3 will not change DS addresses. If this requirement cannot be satisfied (that is, the feature is enabled on a per thread/process basis), then the operating system must ensure that the feature is enabled/disabled appropriately in the context switch code.

### 18.6.4 Performance Monitoring and Intel Hyper-Threading Technology in Processors Based on Intel NetBurst® Microarchitecture

The performance monitoring capability of processors based on Intel NetBurst microarchitecture and supporting Intel Hyper-Threading Technology is similar to that described in Section 18.6.3. However, the capability is extended so that:

- Performance counters can be programmed to select events qualified by logical processor IDs.
- Performance monitoring interrupts can be directed to a specific logical processor within the physical processor.

The sections below describe performance counters, event qualification by logical processor ID, and special purpose bits in ESCRs/CCCRs. They also describe MSR\_PEBS\_ENABLE, MSR\_PEBS\_MATRIX\_VERT, and MSR\_TC\_PRECISE\_EVENT.

#### 18.6.4.1 ESCR MSRs

Figure 18-49 shows the layout of an ESCR MSR in processors supporting Intel Hyper-Threading Technology.

The functions of the flags and fields are as follows:

- **T1\_USR flag, bit 0** — When set, events are counted when thread 1 (logical processor 1) is executing at a current privilege level (CPL) of 1, 2, or 3. These privilege levels are generally used by application code and unprotected operating system code.




Figure 18-49. Event Selection Control Register (ESCR) for the Pentium 4 Processor, Intel Xeon Processor and Intel Xeon Processor MP Supporting Hyper-Threading Technology

- **T1\_OS flag, bit 1** — When set, events are counted when thread 1 (logical processor 1) is executing at CPL of 0. This privilege level is generally reserved for protected operating system code. (When both the T1\_OS and T1\_USR flags are set, thread 1 events are counted at all privilege levels.)
- **T0\_USR flag, bit 2** — When set, events are counted when thread 0 (logical processor 0) is executing at a CPL of 1, 2, or 3.
- **T0\_OS flag, bit 3** — When set, events are counted when thread 0 (logical processor 0) is executing at CPL of 0. (When both the T0\_OS and T0\_USR flags are set, thread 0 events are counted at all privilege levels.)
- **Tag enable, bit 4** — When set, enables tagging of  $\mu$ ops to assist in at-retirement event counting; when clear, disables tagging. See Section 18.6.3.6, “At-Retirement Counting.”
- **Tag value field, bits 5 through 8** — Selects a tag value to associate with a  $\mu$ op to assist in at-retirement event counting.
- **Event mask field, bits 9 through 24** — Selects events to be counted from the event class selected with the event select field.
- **Event select field, bits 25 through 30** — Selects a class of events to be counted. The events within this class that are counted are selected with the event mask field.

The T0\_OS and T0\_USR flags and the T1\_OS and T1\_USR flags allow event counting and sampling to be specified for a specific logical processor (0 or 1) within an Intel Xeon processor MP (See also: Section 8.4.5, “Identifying Logical Processors in an MP System,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*).

Not all performance monitoring events can be detected within an Intel Xeon processor MP on a per logical processor basis (see Section 18.6.4.4, “Performance Monitoring Events”). Some sub-events (specified by an event mask bits) are counted or sampled without regard to which logical processor is associated with the detected event.

#### 18.6.4.2 CCCR MSRs

Figure 18-50 shows the layout of a CCCR MSR in processors supporting Intel Hyper-Threading Technology. The functions of the flags and fields are as follows:

- **Enable flag, bit 12** — When set, enables counting; when clear, the counter is disabled. This flag is cleared on reset
- **ESCR select field, bits 13 through 15** — Identifies the ESCR to be used to select events to be counted with the counter associated with the CCCR.
- **Active thread field, bits 16 and 17** — Enables counting depending on which logical processors are active (executing a thread). This field enables filtering of events based on the state (active or inactive) of the logical processors. The encodings of this field are as follows:
  - 00 — None. Count only when neither logical processor is active.
  - 01 — Single. Count only when one logical processor is active (either 0 or 1).
  - 10 — Both. Count only when both logical processors are active.
  - 11 — Any. Count when either logical processor is active.
 A halted logical processor or a logical processor in the “wait for SIPI” state is considered inactive.
- **Compare flag, bit 18** — When set, enables filtering of the event count; when clear, disables filtering. The filtering method is selected with the threshold, complement, and edge flags.




Figure 18-50. Counter Configuration Control Register (CCCR)

- **Complement flag, bit 19** — Selects how the incoming event count is compared with the threshold value. When set, event counts that are less than or equal to the threshold value result in a single count being delivered to the performance counter; when clear, counts greater than the threshold value result in a count being delivered to the performance counter (see Section 18.6.3.5.2, “Filtering Events”). The compare flag is not active unless the compare flag is set.
- **Threshold field, bits 20 through 23** — Selects the threshold value to be used for comparisons. The processor examines this field only when the compare flag is set, and uses the complement flag setting to determine the type of threshold comparison to be made. The useful range of values that can be entered in this field depend on the type of event being counted (see Section 18.6.3.5.2, “Filtering Events”).
- **Edge flag, bit 24** — When set, enables rising edge (false-to-true) edge detection of the threshold comparison output for filtering event counts; when clear, rising edge detection is disabled. This flag is active only when the compare flag is set.
- **FORCE\_OVF flag, bit 25** — When set, forces a counter overflow on every counter increment; when clear, overflow only occurs when the counter actually overflows.
- **OVF\_PMI\_T0 flag, bit 26** — When set, causes a performance monitor interrupt (PMI) to be sent to logical processor 0 when the counter overflows occurs; when clear, disables PMI generation for logical processor 0. Note that the PMI is generate on the next event count after the counter has overflowed.
- **OVF\_PMI\_T1 flag, bit 27** — When set, causes a performance monitor interrupt (PMI) to be sent to logical processor 1 when the counter overflows occurs; when clear, disables PMI generation for logical processor 1. Note that the PMI is generate on the next event count after the counter has overflowed.
- **Cascade flag, bit 30** — When set, enables counting on one counter of a counter pair when its alternate counter in the other the counter pair in the same counter group overflows (see Section 18.6.3.2, “Performance Counters,” for further details); when clear, disables cascading of counters.
- **OVF flag, bit 31** — Indicates that the counter has overflowed when set. This flag is a sticky flag that must be explicitly cleared by software.

### 18.6.4.3 IA32\_PEBS\_ENABLE MSR

In a processor supporting Intel Hyper-Threading Technology and based on the Intel NetBurst microarchitecture, PEBS is enabled and qualified with two bits in the MSR\_PEBS\_ENABLE MSR: bit 25 (ENABLE\_PEBS\_MY\_THR) and 26 (ENABLE\_PEBS\_OTH\_THR) respectively. These bits do not explicitly identify a specific logical processor by logic

processor ID(T0 or T1); instead, they allow a software agent to enable PEBS for subsequent threads of execution on the same logical processor on which the agent is running (“my thread”) or for the other logical processor in the physical package on which the agent is not running (“other thread”).

PEBS is supported for only a subset of the at-retirement events: Execution\_event, Front\_end\_event, and Replay\_event. Also, PEBS can be carried out only with two performance counters: MSR\_IQ\_CCCR4 (MSR address 370H) for logical processor 0 and MSR\_IQ\_CCCR5 (MSR address 371H) for logical processor 1.

Performance monitoring tools should use a processor affinity mask to bind the kernel mode components that need to modify the ENABLE\_PEBS\_MY\_THR and ENABLE\_PEBS\_OTH\_THR bits in the MSR\_PEBS\_ENABLE MSR to a specific logical processor. This is to prevent these kernel mode components from migrating between different logical processors due to OS scheduling.

#### 18.6.4.4 Performance Monitoring Events

All of the events listed in Table 19-31 and 19-32 are available in an Intel Xeon processor MP. When Intel Hyper-Threading Technology is active, many performance monitoring events can be qualified by the logical processor ID, which corresponds to bit 0 of the initial APIC ID. This allows for counting an event in any or all of the logical processors. However, not all the events have this logic processor specificity, or thread specificity.

Here, each event falls into one of two categories:

- **Thread specific (TS)** — The event can be qualified as occurring on a specific logical processor.
- **Thread independent (TI)** — The event cannot be qualified as being associated with a specific logical processor.

Table 19-37 gives logical processor specific information (TS or TI) for each of the events described in Tables 19-31 and 19-32. If for example, a TS event occurred in logical processor T0, the counting of the event (as shown in Table 18-73) depends only on the setting of the T0\_USR and T0\_OS flags in the ESCR being used to set up the event counter. The T1\_USR and T1\_OS flags have no effect on the count.

**Table 18-73. Effect of Logical Processor and CPL Qualification for Logical-Processor-Specific (TS) Events**

|                   | T1_OS/T1_USR = 00            | T1_OS/T1_USR = 01   | T1_OS/T1_USR = 11   | T1_OS/T1_USR = 10  |
|-------------------|------------------------------|---|---|--|
| T0_OS/T0_USR = 00 | Zero count                   | Counts while T1 in USR                                      | Counts while T1 in OS or USR                                | Counts while T1 in OS                                      |
| T0_OS/T0_USR = 01 | Counts while T0 in USR       | Counts while T0 in USR or T1 in USR                         | Counts while (a) T0 in USR or (b) T1 in OS or (c) T1 in USR | Counts while (a) T0 in OS or (b) T1 in OS                  |
| T0_OS/T0_USR = 11 | Counts while T0 in OS or USR | Counts while (a) T0 in OS or (b) T0 in USR or (c) T1 in USR | Counts irrespective of CPL, T0, T1                          | Counts while (a) T0 in OS or (b) T0 in USR or (c) T1 in OS |
| T0_OS/T0_USR = 10 | Counts T0 in OS              | Counts T0 in OS or T1 in USR                                | Counts while (a) T0 in OS or (b) T1 in OS or (c) T1 in USR  | Counts while (a) T0 in OS or (b) T1 in OS                  |

When a bit in the event mask field is TI, the effect of specifying bit-0-3 of the associated ESCR are described in Table 15-6. For events that are marked as TI in Chapter 19, the effect of selectively specifying T0\_USR, T0\_OS, T1\_USR, T1\_OS bits is shown in Table 18-74.

**Table 18-74. Effect of Logical Processor and CPL Qualification for Non-logical-Processor-specific (TI) Events**

|                   | T1_OS/T1_USR = 00                           | T1_OS/T1_USR = 01                           | T1_OS/T1_USR = 11                  | T1_OS/T1_USR = 10                         |
|-------------------|---|---|------------------------------------|---|
| T0_OS/T0_USR = 00 | Zero count                                  | Counts while (a) T0 in USR or (b) T1 in USR | Counts irrespective of CPL, T0, T1 | Counts while (a) T0 in OS or (b) T1 in OS |
| T0_OS/T0_USR = 01 | Counts while (a) T0 in USR or (b) T1 in USR | Counts while (a) T0 in USR or (b) T1 in USR | Counts irrespective of CPL, T0, T1 | Counts irrespective of CPL, T0, T1        |
| T0_OS/T0_USR = 11 | Counts irrespective of CPL, T0, T1          | Counts irrespective of CPL, T0, T1          | Counts irrespective of CPL, T0, T1 | Counts irrespective of CPL, T0, T1        |
| T0_OS/T0_USR = 0  | Counts while (a) T0 in OS or (b) T1 in OS   | Counts irrespective of CPL, T0, T1          | Counts irrespective of CPL, T0, T1 | Counts while (a) T0 in OS or (b) T1 in OS |

### 18.6.4.5 Counting Clocks on systems with Intel Hyper-Threading Technology in Processors Based on Intel NetBurst® Microarchitecture

#### 18.6.4.5.1 Non-Halted Clockticks

Use the following procedure to program ESCRs and CCCRs to obtain non-halted clockticks on processors based on Intel NetBurst microarchitecture:

1. Select an ESCR for the `global_power_events` and specify the `RUNNING` sub-event mask and the desired `T0_OS/T0_USR/T1_OS/T1_USR` bits for the targeted processor.
2. Select an appropriate counter.
3. Enable counting in the CCCR for that counter by setting the enable bit.

#### 18.6.4.5.2 Non-Sleep Clockticks

Performance monitoring counters can be configured to count clockticks whenever the performance monitoring hardware is not powered-down. To count Non-sleep Clockticks with a performance-monitoring counter, do the following:

1. Select one of the 18 counters.
2. Select any of the ESCRs whose events the selected counter can count. Set its event select to anything other than "no\_event"; the counter may be disabled if this is not done.
3. Turn threshold comparison on in the CCCR by setting the compare bit to "1".
4. Set the threshold to "15" and the complement to "1" in the CCCR. Since no event can exceed this threshold, the threshold condition is met every cycle and the counter counts every cycle. Note that this overrides any qualification (e.g. by CPL) specified in the ESCR.
5. Enable counting in the CCCR for the counter by setting the enable bit.

In most cases, the counts produced by the non-halted and non-sleep metrics are equivalent if the physical package supports one logical processor and is not placed in a power-saving state. Operating systems may execute an `HLT` instruction and place a physical processor in a power-saving state.

On processors that support Intel Hyper-Threading Technology (Intel HT Technology), each physical package can support two or more logical processors. Current implementation of Intel HT Technology provides two logical processors for each physical processor. While both logical processors can execute two threads simultaneously, one logical processor may halt to allow the other logical processor to execute without sharing execution resources between two logical processors.

Non-halted Clockticks can be set up to count the number of processor clock cycles for each logical processor whenever the logical processor is not halted (the count may include some portion of the clock cycles for that logical processor to complete a transition to a halted state). Physical processors that support Intel HT Technology enter into a power-saving state if all logical processors halt.



The Non-sleep Clockticks mechanism uses a filtering mechanism in CCCRs. The mechanism will continue to increment as long as one logical processor is not halted or in a power-saving state. Applications may cause a processor to enter into a power-saving state by using an OS service that transfers control to an OS's idle loop. The idle loop then may place the processor into a power-saving state after an implementation-dependent period if there is no work for the processor.

### 18.6.5 Performance Monitoring and Dual-Core Technology

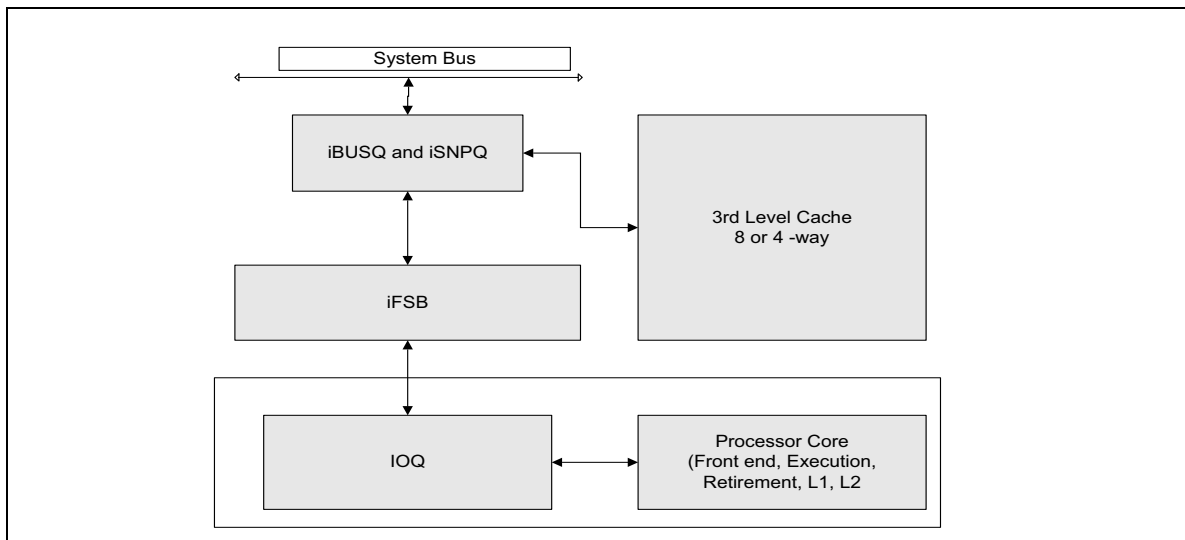
The performance monitoring capability of dual-core processors duplicates the microarchitectural resources of a single-core processor implementation. Each processor core has dedicated performance monitoring resources.

In the case of Pentium D processor, each logical processor is associated with dedicated resources for performance monitoring. In the case of Pentium processor Extreme edition, each processor core has dedicated resources, but two logical processors in the same core share performance monitoring resources (see Section 18.6.4, "Performance Monitoring and Intel Hyper-Threading Technology in Processors Based on Intel NetBurst<sup>®</sup> Microarchitecture").

### 18.6.6 Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache

The 64-bit Intel Xeon processor MP with up to 8-MByte L3 cache has a CPUID signature of family [0FH], model [03H or 04H]. Performance monitoring capabilities available to Pentium 4 and Intel Xeon processors with the same values (see Section 18.1 and Section 18.6.4) apply to the 64-bit Intel Xeon processor MP with an L3 cache.

The level 3 cache is connected between the system bus and IOQ through additional control logic. See Figure 18-51.



**Figure 18-51. Block Diagram of 64-bit Intel Xeon Processor MP with 8-MByte L3**

Additional performance monitoring capabilities and facilities unique to 64-bit Intel Xeon processor MP with an L3 cache are described in this section. The facility for monitoring events consists of a set of dedicated model-specific registers (MSRs), each dedicated to a specific event. Programming of these MSRs requires using RDMSR/WRMSR instructions with 64-bit values.

The lower 32-bits of the MSRs at addresses 107CC through 107D3 are treated as 32 bit performance counter registers. These performance counters can be accessed using RDPMC instruction with the index starting from 18 through 25. The EDX register returns zero when reading these 8 PMCs.

The performance monitoring capabilities consist of four events. These are:

- IBUSQ event** — This event detects the occurrence of micro-architectural conditions related to the iBUSQ unit. It provides two MSRs: MSR\_IFSB\_IBUSQ0 and MSR\_IFSB\_IBUSQ1. Configure sub-event qualification and enable/disable functions using the high 32 bits of these MSRs. The low 32 bits act as a 32-bit event counter. Counting starts after software writes a non-zero value to one or more of the upper 32 bits. See Figure 18-52.




Figure 18-52. MSR\_IFSB\_IBUSQx, Addresses: 107CCH and 107CDH

- ISNPQ event** — This event detects the occurrence of microarchitectural conditions related to the iSNPQ unit. It provides two MSRs: MSR\_IFSB\_ISNPQ0 and MSR\_IFSB\_ISNPQ1. Configure sub-event qualifications and enable/disable functions using the high 32 bits of the MSRs. The low 32-bits act as a 32-bit event counter. Counting starts after software writes a non-zero value to one or more of the upper 32-bits. See Figure 18-53.




Figure 18-53. MSR\_IFSB\_ISNPQx, Addresses: 107CEH and 107CFH

- EFSB event** — This event can detect the occurrence of micro-architectural conditions related to the iFSB unit or system bus. It provides two MSRs: MSR\_EFSB\_DRDY0 and MSR\_EFSB\_DRDY1. Configure sub-event qualifications and enable/disable functions using the high 32 bits of the 64-bit MSR. The low 32-bit act as a 32-bit event counter. Counting starts after software writes a non-zero value to one or more of the qualification bits in the upper 32-bits of the MSR. See Figure 18-54.




Figure 18-54. MSR\_EFSB\_DRDYx, Addresses: 107D0H and 107D1H

- iBUSQ Latency event** — This event accumulates weighted cycle counts for latency measurement of transactions in the iBUSQ unit. The count is enabled by setting MSR\_IFSB\_CTRL6[bit 26] to 1; the count freezes after software sets MSR\_IFSB\_CTRL6[bit 26] to 0. MSR\_IFSB\_CNTR7 acts as a 64-bit event counter for this event. See Figure 18-55.

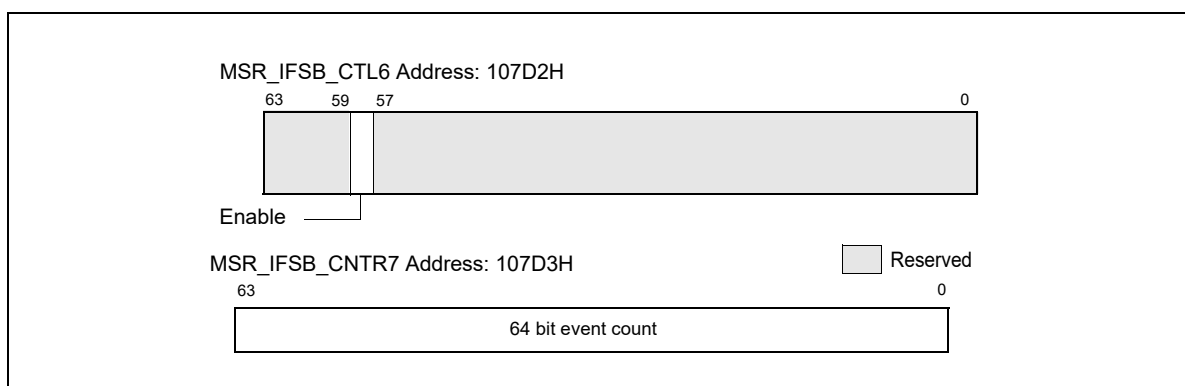


Figure 18-55. MSR\_IFSB\_CTL6, Address: 107D2H;  
 MSR\_IFSB\_CNTR7, Address: 107D3H

## 18.6.7 Performance Monitoring on L3 and Caching Bus Controller Sub-Systems


The Intel Xeon processor 7400 series and Dual-Core Intel Xeon processor 7100 series employ a distinct L3/caching bus controller sub-system. These sub-system have a unique set of performance monitoring capability and programming interfaces that are largely common between these two processor families.

Intel Xeon processor 7400 series are based on 45 nm enhanced Intel Core microarchitecture. The CPUID signature is indicated by DisplayFamily\_DisplayModel value of 06\_1DH (see CPUID instruction in Chapter 3, "Instruction Set Reference, A-L" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A*). Intel Xeon processor 7400 series have six processor cores that share an L3 cache.

Dual-Core Intel Xeon processor 7100 series are based on Intel NetBurst microarchitecture, have a CPUID signature of family [0FH], model [06H] and a unified L3 cache shared between two cores. Each core in an Intel Xeon processor 7100 series supports Intel Hyper-Threading Technology, providing two logical processors per core.

Both Intel Xeon processor 7400 series and Intel Xeon processor 7100 series support multi-processor configurations using system bus interfaces. In Intel Xeon processor 7400 series, the L3/caching bus controller sub-system provides three Simple Direct Interface (SDI) to service transactions originated the XQ-replacement SDI logic in each dual-core modules. In Intel Xeon processor 7100 series, the IOQ logic in each processor core is replaced with a Simple Direct Interface (SDI) logic. The L3 cache is connected between the system bus and the SDI through

additional control logic. See Figure 18-56 for the block configuration of six processor cores and the L3/Caching bus controller sub-system in Intel Xeon processor 7400 series. Figure 18-56 shows the block configuration of two processor cores (four logical processors) and the L3/Caching bus controller sub-system in Intel Xeon processor 7100 series.



**Figure 18-56. Block Diagram of Intel Xeon Processor 7400 Series**

Almost all of the performance monitoring capabilities available to processor cores with the same CPUID signatures (see Section 18.1 and Section 18.6.4) apply to Intel Xeon processor 7100 series. The MSR's used by performance monitoring interface are shared between two logical processors in the same processor core.

The performance monitoring capabilities available to processor with DisplayFamily\_DisplayModel signature 06\_17H also apply to Intel Xeon processor 7400 series. Each processor core provides its own set of MSR's for performance monitoring interface.

The IOQ\_allocation and IOQ\_active\_entries events are not supported in Intel Xeon processor 7100 series and 7400 series. Additional performance monitoring capabilities applicable to the L3/caching bus controller sub-system are described in this section.




Figure 18-57. Block Diagram of Intel Xeon Processor 7100 Series

### 18.6.7.1 Overview of Performance Monitoring with L3/Caching Bus Controller

The facility for monitoring events consists of a set of dedicated model-specific registers (MSRs). There are eight event select/counting MSRs that are dedicated to counting events associated with specified microarchitectural conditions. Programming of these MSRs requires using RDMSR/WRMSR instructions with 64-bit values. In addition, an MSR MSR\_EMON\_L3\_GL\_CTL provides simplified interface to control freezing, resetting, re-enabling operation of any combination of these event select/counting MSRs.

The eight MSRs dedicated to count occurrences of specific conditions are further divided to count three sub-classes of microarchitectural conditions:

- Two MSRs (MSR\_EMON\_L3\_CTR\_CTL0 and MSR\_EMON\_L3\_CTR\_CTL1) are dedicated to counting GBSQ events. Up to two GBSQ events can be programmed and counted simultaneously.
- Two MSRs (MSR\_EMON\_L3\_CTR\_CTL2 and MSR\_EMON\_L3\_CTR\_CTL3) are dedicated to counting GSNPQ events. Up to two GSNPQ events can be programmed and counted simultaneously.
- Four MSRs (MSR\_EMON\_L3\_CTR\_CTL4, MSR\_EMON\_L3\_CTR\_CTL5, MSR\_EMON\_L3\_CTR\_CTL6, and MSR\_EMON\_L3\_CTR\_CTL7) are dedicated to counting external bus operations.

The bit fields in each of eight MSRs share the following common characteristics:

- Bits 63:32 is the event control field that includes an event mask and other bit fields that control counter operation. The event mask field specifies details of the microarchitectural condition, and its definition differs across GBSQ, GSNPQ, FSB.
- Bits 31:0 is the event count field. If the specified condition is met during each relevant clock domain of the event logic, the matched condition signals the counter logic to increment the associated event count field. The lower 32-bits of these 8 MSRs at addresses 107CC through 107D3 are treated as 32 bit performance counter registers.

In Dual-Core Intel Xeon processor 7100 series, the uncore performance counters can be accessed using RDPMC instruction with the index starting from 18 through 25. The EDX register returns zero when reading these 8 PMCs.

In Intel Xeon processor 7400 series, RDPMC with ECX between 2 and 9 can be used to access the eight uncore performance counter/control registers.

### 18.6.7.2 GBSQ Event Interface

The layout of MSR\_EMON\_L3\_CTR\_CTL0 and MSR\_EMON\_L3\_CTR\_CTL1 is given in Figure 18-58. Counting starts after software writes a non-zero value to one or more of the upper 32 bits.

The event mask field (bits 58:32) consists of the following eight attributes:

- Agent\_Select (bits 35:32): The definition of this field differs slightly between Intel Xeon processor 7100 and 7400.

For Intel Xeon processor 7100 series, each bit specifies a logical processor in the physical package. The lower two bits corresponds to two logical processors in the first processor core, the upper two bits corresponds to two logical processors in the second processor core. 0FH encoding matches transactions from any logical processor.

For Intel Xeon processor 7400 series, each bit of [34:32] specifies the SDI logic of a dual-core module as the originator of the transaction. A value of 0111B in bits [35:32] specifies transaction from any processor core.

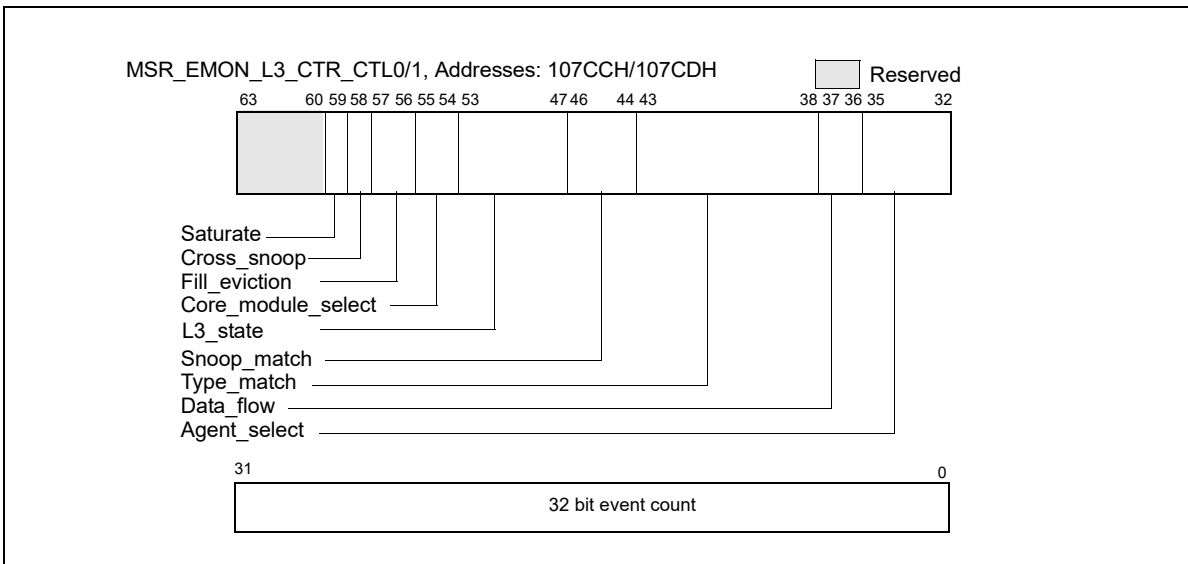


Figure 18-58. MSR\_EMON\_L3\_CTR\_CTL0/1, Addresses: 107CCH/107CDH

- Data\_Flow (bits 37:36): Bit 36 specifies demand transactions, bit 37 specifies prefetch transactions.
- Type\_Match (bits 43:38): Specifies transaction types. If all six bits are set, event count will include all transaction types.
- Snoop\_Match: (bits 46:44): The three bits specify (in ascending bit position) clean snoop result, HIT snoop result, and HITM snoop results respectively.
- L3\_State (bits 53:47): Each bit specifies an L2 coherency state.
- Core\_Module\_Select (bits 55:54): The valid encodings for L3 lookup differ slightly between Intel Xeon processor 7100 and 7400.

For Intel Xeon processor 7100 series,

- 00B: Match transactions from any core in the physical package
- 01B: Match transactions from this core only
- 10B: Match transactions from the other core in the physical package
- 11B: Match transaction from both cores in the physical package

For Intel Xeon processor 7400 series,

- 00B: Match transactions from any dual-core module in the physical package
- 01B: Match transactions from this dual-core module only
- 10B: Match transactions from either one of the other two dual-core modules in the physical package

- 11B: Match transaction from more than one dual-core modules in the physical package
- Fill\_Eviction (bits 57:56): The valid encodings are
  - 00B: Match any transactions
  - 01B: Match transactions that fill L3
  - 10B: Match transactions that fill L3 without an eviction
  - 11B: Match transaction fill L3 with an eviction
- Cross\_Snoop (bit 58): The encodings are
  - 0B: Match any transactions
  - 1B: Match cross snoop transactions

For each counting clock domain, if all eight attributes match, event logic signals to increment the event count field.

### 18.6.7.3 GSNPQ Event Interface

The layout of MSR\_EMON\_L3\_CTR\_CTL2 and MSR\_EMON\_L3\_CTR\_CTL3 is given in Figure 18-59. Counting starts after software writes a non-zero value to one or more of the upper 32 bits.

The event mask field (bits 58:32) consists of the following six attributes:

- Agent\_Select (bits 37:32): The definition of this field differs slightly between Intel Xeon processor 7100 and 7400.
- For Intel Xeon processor 7100 series, each of the lowest 4 bits specifies a logical processor in the physical package. The lowest two bits corresponds to two logical processors in the first processor core, the next two bits corresponds to two logical processors in the second processor core. Bit 36 specifies other symmetric agent transactions. Bit 37 specifies central agent transactions. 3FH encoding matches transactions from any logical processor.

For Intel Xeon processor 7400 series, each of the lowest 3 bits specifies a dual-core module in the physical package. Bit 37 specifies central agent transactions.

- Type\_Match (bits 43:38): Specifies transaction types. If all six bits are set, event count will include any transaction types.
- Snoop\_Match: (bits 46:44): The three bits specify (in ascending bit position) clean snoop result, HIT snoop result, and HITM snoop results respectively.
- L2\_State (bits 53:47): Each bit specifies an L3 coherency state.
- Core\_Module\_Select (bits 56:54): Bit 56 enables Core\_Module\_Select matching. If bit 56 is clear, Core\_Module\_Select encoding is ignored. The valid encodings for the lower two bits (bit 55, 54) differ slightly between Intel Xeon processor 7100 and 7400.

For Intel Xeon processor 7100 series, if bit 56 is set, the valid encodings for the lower two bits (bit 55, 54) are

- 00B: Match transactions from only one core (irrespective which core) in the physical package
- 01B: Match transactions from this core and not the other core
- 10B: Match transactions from the other core in the physical package, but not this core
- 11B: Match transaction from both cores in the physical package

For Intel Xeon processor 7400 series, if bit 56 is set, the valid encodings for the lower two bits (bit 55, 54) are

- 00B: Match transactions from only one dual-core module (irrespective which module) in the physical package.
- 01B: Match transactions from one or more dual-core modules.
- 10B: Match transactions from two or more dual-core modules.
- 11B: Match transaction from all three dual-core modules in the physical package.

- Block\_Snoop (bit 57): specifies blocked snoop.

For each counting clock domain, if all six attributes match, event logic signals to increment the event count field.




Figure 18-59. MSR\_EMON\_L3\_CTR\_CTL2/3, Addresses: 107CEH/107CFH

### 18.6.7.4 FSB Event Interface

The layout of MSR\_EMON\_L3\_CTR\_CTL4 through MSR\_EMON\_L3\_CTR\_CTL7 is given in Figure 18-60. Counting starts after software writes a non-zero value to one or more of the upper 32 bits.

The event mask field (bits 58:32) is organized as follows:

- Bit 58: must set to 1.
- FSB\_Submask (bits 57:32): Specifies FSB-specific sub-event mask.

The FSB sub-event mask defines a set of independent attributes. The event logic signals to increment the associated event count field if one of the attribute matches. Some of the sub-event mask bit counts durations. A duration event increments at most once per cycle.

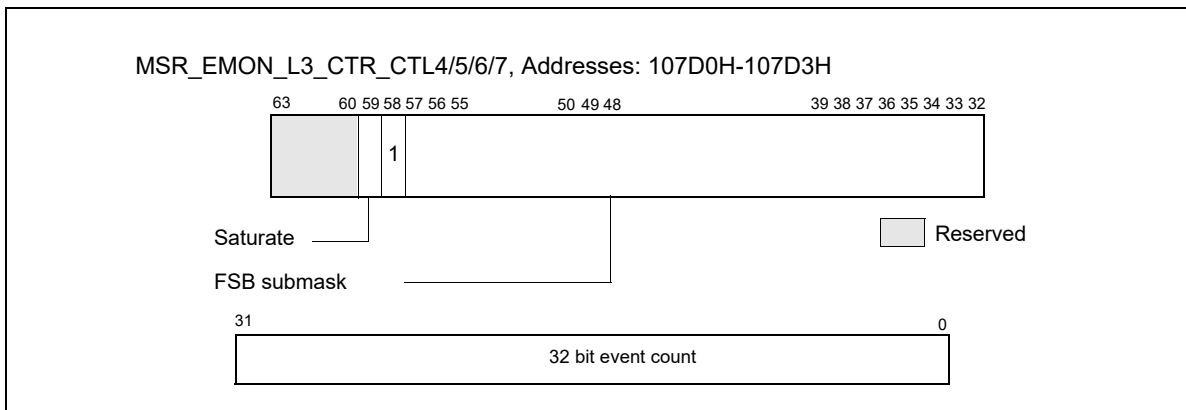


Figure 18-60. MSR\_EMON\_L3\_CTR\_CTL4/5/6/7, Addresses: 107D0H-107D3H

#### 18.6.7.4.1 FSB Sub-Event Mask Interface

- FSB\_type (bit 37:32): Specifies different FSB transaction types originated from this physical package.
- FSB\_L\_clear (bit 38): Count clean snoop results from any source for transaction originated from this physical package.
- FSB\_L\_hit (bit 39): Count HIT snoop results from any source for transaction originated from this physical package.



- FSB\_L\_hitm (bit 40): Count HITM snoop results from any source for transaction originated from this physical package.
- FSB\_L\_defer (bit 41): Count DEFER responses to this processor's transactions.
- FSB\_L\_retry (bit 42): Count RETRY responses to this processor's transactions.
- FSB\_L\_snoop\_stall (bit 43): Count snoop stalls to this processor's transactions.
- FSB\_DBSY (bit 44): Count DBSY assertions by this processor (without a concurrent DRDY).
- FSB\_DRDY (bit 45): Count DRDY assertions by this processor.
- FSB\_BNR (bit 46): Count BNR assertions by this processor.
- FSB\_IOQ\_empty (bit 47): Counts each bus clocks when the IOQ is empty.
- FSB\_IOQ\_full (bit 48): Counts each bus clocks when the IOQ is full.
- FSB\_IOQ\_active (bit 49): Counts each bus clocks when there is at least one entry in the IOQ.
- FSB\_WW\_data (bit 50): Counts back-to-back write transaction's data phase.
- FSB\_WW\_issue (bit 51): Counts back-to-back write transaction request pairs issued by this processor.
- FSB\_WR\_issue (bit 52): Counts back-to-back write-read transaction request pairs issued by this processor.
- FSB\_RW\_issue (bit 53): Counts back-to-back read-write transaction request pairs issued by this processor.
- FSB\_other\_DBSY (bit 54): Count DBSY assertions by another agent (without a concurrent DRDY).
- FSB\_other\_DRDY (bit 55): Count DRDY assertions by another agent.
- FSB\_other\_snoop\_stall (bit 56): Count snoop stalls on the FSB due to another agent.
- FSB\_other\_BNR (bit 57): Count BNR assertions from another agent.

#### 18.6.7.5 Common Event Control Interface

The MSR\_EMON\_L3\_GL\_CTL MSR provides simplified access to query overflow status of the GBSQ, GSNPQ, FSB event counters. It also provides control bit fields to freeze, unfreeze, or reset those counters. The following bit fields are supported:

- GL\_freeze\_cmd (bit 0): Freeze the event counters specified by the GL\_event\_select field.
- GL\_unfreeze\_cmd (bit 1): Unfreeze the event counters specified by the GL\_event\_select field.
- GL\_reset\_cmd (bit 2): Clear the event count field of the event counters specified by the GL\_event\_select field. The event select field is not affected.
- GL\_event\_select (bit 23:16): Selects one or more event counters to subject to specified command operations indicated by bits 2:0. Bit 16 corresponds to MSR\_EMON\_L3\_CTR\_CTL0, bit 23 corresponds to MSR\_EMON\_L3\_CTR\_CTL7.
- GL\_event\_status (bit 55:48): Indicates the overflow status of each event counters. Bit 48 corresponds to MSR\_EMON\_L3\_CTR\_CTL0, bit 55 corresponds to MSR\_EMON\_L3\_CTR\_CTL7.

In the event control field (bits 63:32) of each MSR, if the saturate control (bit 59, see Figure 18-58 for example) is set, the event logic forces the value FFFF\_FFFFH into the event count field instead of incrementing it.

#### 18.6.8 Performance Monitoring (P6 Family Processor)

The P6 family processors provide two 40-bit performance counters, allowing two types of events to be monitored simultaneously. These can either count events or measure duration. When counting events, a counter increments each time a specified event takes place or a specified number of events takes place. When measuring duration, it counts the number of processor clocks that occur while a specified condition is true. The counters can count events or measure durations that occur at any privilege level.

Table 19-40, Chapter 19, lists the events that can be counted with the P6 family performance monitoring counters.

**NOTE**

The performance-monitoring events listed in Chapter 19 are intended to be used as guides for performance tuning. Counter values reported are not guaranteed to be accurate and should be used as a relative guide for tuning. Known discrepancies are documented where applicable.

The performance-monitoring counters are supported by four MSR: the performance event select MSRs (PerfEvtSel0 and PerfEvtSel1) and the performance counter MSRs (PerfCtr0 and PerfCtr1). These registers can be read from and written to using the RDMSR and WRMSR instructions, respectively. They can be accessed using these instructions only when operating at privilege level 0. The PerfCtr0 and PerfCtr1 MSRs can be read from any privilege level using the RDPNC (read performance-monitoring counters) instruction.

**NOTE**


The PerfEvtSel0, PerfEvtSel1, PerfCtr0, and PerfCtr1 MSRs and the events listed in Table 19-40 are model-specific for P6 family processors. They are not guaranteed to be available in other IA-32 processors.

**18.6.8.1 PerfEvtSel0 and PerfEvtSel1 MSRs**

The PerfEvtSel0 and PerfEvtSel1 MSRs control the operation of the performance-monitoring counters, with one register used to set up each counter. They specify the events to be counted, how they should be counted, and the privilege levels at which counting should take place. Figure 18-61 shows the flags and fields in these MSRs.

The functions of the flags and fields in the PerfEvtSel0 and PerfEvtSel1 MSRs are as follows:

- **Event select field (bits 0 through 7)** — Selects the event logic unit to detect certain microarchitectural conditions (see Table 19-40, for a list of events and their 8-bit codes).
- **Unit mask (UMASK) field (bits 8 through 15)** — Further qualifies the event logic unit selected in the event select field to detect a specific microarchitectural condition. For example, for some cache events, the mask is used as a MESI-protocol qualifier of cache states (see Table 19-40).



**Figure 18-61. PerfEvtSel0 and PerfEvtSel1 MSRs**

- **USR (user mode) flag (bit 16)** — Specifies that events are counted only when the processor is operating at privilege levels 1, 2 or 3. This flag can be used in conjunction with the OS flag.
- **OS (operating system mode) flag (bit 17)** — Specifies that events are counted only when the processor is operating at privilege level 0. This flag can be used in conjunction with the USR flag.
- **E (edge detect) flag (bit 18)** — Enables (when set) edge detection of events. The processor counts the number of deasserted to asserted transitions of any condition that can be expressed by the other fields. The mechanism is limited in that it does not permit back-to-back assertions to be distinguished. This mechanism allows software to measure not only the fraction of time spent in a particular state, but also the average length of time spent in such a state (for example, the time spent waiting for an interrupt to be serviced).

- **PC (pin control) flag (bit 19)** — When set, the processor toggles the PM*i* pins and increments the counter when performance-monitoring events occur; when clear, the processor toggles the PM*i* pins when the counter overflows. The toggling of a pin is defined as assertion of the pin for a single bus clock followed by deassertion.
- **INT (APIC interrupt enable) flag (bit 20)** — When set, the processor generates an exception through its local APIC on counter overflow.
- **EN (Enable Counters) Flag (bit 22)** — This flag is only present in the PerfEvtSel0 MSR. When set, performance counting is enabled in both performance-monitoring counters; when clear, both counters are disabled.
- **INV (invert) flag (bit 23)** — When set, inverts the counter-mask (CMASK) comparison, so that both greater than or equal to and less than comparisons can be made (0: greater than or equal; 1: less than). Note if counter-mask is programmed to zero, INV flag is ignored.
- **Counter mask (CMASK) field (bits 24 through 31)** — When nonzero, the processor compares this mask to the number of events counted during a single cycle. If the event count is greater than or equal to this mask, the counter is incremented by one. Otherwise the counter is not incremented. This mask can be used to count events only if multiple occurrences happen per clock (for example, two or more instructions retired per clock). If the counter-mask field is 0, then the counter is incremented each cycle by the number of events that occurred that cycle.

### 18.6.8.2 PerfCtr0 and PerfCtr1 MSRs

The performance-counter MSRs (PerfCtr0 and PerfCtr1) contain the event or duration counts for the selected events being counted. The RDPMC instruction can be used by programs or procedures running at any privilege level and in virtual-8086 mode to read these counters. The PCE flag in control register CR4 (bit 8) allows the use of this instruction to be restricted to only programs and procedures running at privilege level 0.

The RDPMC instruction is not serializing or ordered with other instructions. Thus, it does not necessarily wait until all previous instructions have been executed before reading the counter. Similarly, subsequent instructions may begin execution before the RDPMC instruction operation is performed.

Only the operating system, executing at privilege level 0, can directly manipulate the performance counters, using the RDMSR and WRMSR instructions. A secure operating system would clear the PCE flag during system initialization to disable direct user access to the performance-monitoring counters, but provide a user-accessible programming interface that emulates the RDPMC instruction.

The WRMSR instruction cannot arbitrarily write to the performance-monitoring counter MSRs (PerfCtr0 and PerfCtr1). Instead, the lower-order 32 bits of each MSR may be written with any value, and the high-order 8 bits are sign-extended according to the value of bit 31. This operation allows writing both positive and negative values to the performance counters.

### 18.6.8.3 Starting and Stopping the Performance-Monitoring Counters

The performance-monitoring counters are started by writing valid setup information in the PerfEvtSel0 and/or PerfEvtSel1 MSRs and setting the enable counters flag in the PerfEvtSel0 MSR. If the setup is valid, the counters begin counting following the execution of a WRMSR instruction that sets the enable counter flag. The counters can be stopped by clearing the enable counters flag or by clearing all the bits in the PerfEvtSel0 and PerfEvtSel1 MSRs. Counter 1 alone can be stopped by clearing the PerfEvtSel1 MSR.

### 18.6.8.4 Event and Time-Stamp Monitoring Software

To use the performance-monitoring counters and time-stamp counter, the operating system needs to provide an event-monitoring device driver. This driver should include procedures for handling the following operations:

- Feature checking.
- Initialize and start counters.
- Stop counters.
- Read the event counters.
- Read the time-stamp counter.

The event monitor feature determination procedure must check whether the current processor supports the performance-monitoring counters and time-stamp counter. This procedure compares the family and model of the processor returned by the CPUID instruction with those of processors known to support performance monitoring. (The Pentium and P6 family processors support performance counters.) The procedure also checks the MSR and TSC flags returned to register EDX by the CPUID instruction to determine if the MSRs and the RDTSC instruction are supported.

The initialize and start counters procedure sets the PerfEvtSel0 and/or PerfEvtSel1 MSRs for the events to be counted and the method used to count them and initializes the counter MSRs (PerfCtr0 and PerfCtr1) to starting counts. The stop counters procedure stops the performance counters (see Section 18.6.8.3, "Starting and Stopping the Performance-Monitoring Counters").

The read counters procedure reads the values in the PerfCtr0 and PerfCtr1 MSRs, and a read time-stamp counter procedure reads the time-stamp counter. These procedures would be provided in lieu of enabling the RDTSC and RDPMC instructions that allow application code to read the counters.

### 18.6.8.5 Monitoring Counter Overflow

The P6 family processors provide the option of generating a local APIC interrupt when a performance-monitoring counter overflows. This mechanism is enabled by setting the interrupt enable flag in either the PerfEvtSel0 or the PerfEvtSel1 MSR. The primary use of this option is for statistical performance sampling.

To use this option, the operating system should do the following things on the processor for which performance events are required to be monitored:

- Provide an interrupt vector for handling the counter-overflow interrupt.
- Initialize the APIC PERF local vector entry to enable handling of performance-monitor counter overflow events.
- Provide an entry in the IDT that points to a stub exception handler that returns without executing any instructions.
- Provide an event monitor driver that provides the actual interrupt handler and modifies the reserved IDT entry to point to its interrupt routine.

When interrupted by a counter overflow, the interrupt handler needs to perform the following actions:

- Save the instruction pointer (EIP register), code-segment selector, TSS segment selector, counter values and other relevant information at the time of the interrupt.
- Reset the counter to its initial setting and return from the interrupt.

An event monitor application utility or another application program can read the information collected for analysis of the performance of the profiled application.

### 18.6.9 Performance Monitoring (Pentium Processors)

The Pentium processor provides two 40-bit performance counters, which can be used to count events or measure duration. The counters are supported by three MSRs: the control and event select MSR (CESR) and the performance counter MSRs (CTR0 and CTR1). These can be read from and written to using the RDMSR and WRMSR instructions, respectively. They can be accessed using these instructions only when operating at privilege level 0.

Each counter has an associated external pin (PM0/BP0 and PM1/BP1), which can be used to indicate the state of the counter to external hardware.

## NOTES

The CESR, CTR0, and CTR1 MSRs and the events listed in Table 19-41 are model-specific for the Pentium processor.

The performance-monitoring events listed in Chapter 19 are intended to be used as guides for performance tuning. Counter values reported are not guaranteed to be accurate and should be used as a relative guide for tuning. Known discrepancies are documented where applicable.

### 18.6.9.1 Control and Event Select Register (CESR)

The 32-bit control and event select MSR (CESR) controls the operation of performance-monitoring counters CTR0 and CTR1 and the associated pins (see Figure 18-62). To control each counter, the CESR register contains a 6-bit event select field (ES0 and ES1), a pin control flag (PC0 and PC1), and a 3-bit counter control field (CC0 and CC1). The functions of these fields are as follows:

- **ES0 and ES1 (event select) fields (bits 0-5, bits 16-21)** — Selects (by entering an event code in the field) up to two events to be monitored. See Table 19-41 for a list of available event codes.




Figure 18-62. CESR MSR (Pentium Processor Only)

- **CC0 and CC1 (counter control) fields (bits 6-8, bits 22-24)** — Controls the operation of the counter. Control codes are as follows:

- 000 — Count nothing (counter disabled).
- 001 — Count the selected event while CPL is 0, 1, or 2.
- 010 — Count the selected event while CPL is 3.
- 011 — Count the selected event regardless of CPL.
- 100 — Count nothing (counter disabled).
- 101 — Count clocks (duration) while CPL is 0, 1, or 2.
- 110 — Count clocks (duration) while CPL is 3.
- 111 — Count clocks (duration) regardless of CPL.

The highest order bit selects between counting events and counting clocks (duration); the middle bit enables counting when the CPL is 3; and the low-order bit enables counting when the CPL is 0, 1, or 2.

- **PC0 and PC1 (pin control) flags (bits 9, 25)** — Selects the function of the external performance-monitoring counter pin (PM0/BP0 and PM1/BP1). Setting one of these flags to 1 causes the processor to assert its associated pin when the counter has overflowed; setting the flag to 0 causes the pin to be asserted when the counter has been incremented. These flags permit the pins to be individually programmed to indicate the overflow or incremented condition. The external signalling of the event on the pins will lag the internal event by a few clocks as the signals are latched and buffered.

While a counter need not be stopped to sample its contents, it must be stopped and cleared or preset before switching to a new event. It is not possible to set one counter separately. If only one event needs to be changed, the CESR register must be read, the appropriate bits modified, and all bits must then be written back to CESR. At reset, all bits in the CESR register are cleared.

### 18.6.9.2 Use of the Performance-Monitoring Pins

When performance-monitor pins PM0/BP0 and/or PM1/BP1 are configured to indicate when the performance-monitor counter has incremented and an "occurrence event" is being counted, the associated pin is asserted (high) each time the event occurs. When a "duration event" is being counted, the associated PM pin is asserted for the

entire duration of the event. When the performance-monitor pins are configured to indicate when the counter has overflowed, the associated PM pin is asserted when the counter has overflowed.

When the PM0/BP0 and/or PM1/BP1 pins are configured to signal that a counter has incremented, it should be noted that although the counters may increment by 1 or 2 in a single clock, the pins can only indicate that the event occurred. Moreover, since the internal clock frequency may be higher than the external clock frequency, a single external clock may correspond to multiple internal clocks.

A “count up to” function may be provided when the event pin is programmed to signal an overflow of the counter. Because the counters are 40 bits, a carry out of bit 39 indicates an overflow. A counter may be preset to a specific value less than  $2^{40} - 1$ . After the counter has been enabled and the prescribed number of events has transpired, the counter will overflow.

Approximately 5 clocks later, the overflow is indicated externally and appropriate action, such as signaling an interrupt, may then be taken.

The PM0/BP0 and PM1/BP1 pins also serve to indicate breakpoint matches during in-circuit emulation, during which time the counter increment or overflow function of these pins is not available. After RESET, the PM0/BP0 and PM1/BP1 pins are configured for performance monitoring, however a hardware debugger may reconfigure these pins to indicate breakpoint matches.

### 18.6.9.3 Events Counted

Events that performance-monitoring counters can be set to count and record (using CTR0 and CTR1) are divided in two categories: occurrence and duration:

- **Occurrence events** — Counts are incremented each time an event takes place. If PM0/BP0 or PM1/BP1 pins are used to indicate when a counter increments, the pins are asserted each clock counters increment. But if an event happens twice in one clock, the counter increments by 2 (the pins are asserted only once).
- **Duration events** — Counters increment the total number of clocks that the condition is true. When used to indicate when counters increment, PM0/BP0 and/or PM1/BP1 pins are asserted for the duration.

## 18.7 COUNTING CLOCKS

The count of cycles, also known as clockticks, forms the basis for measuring how long a program takes to execute. Clockticks are also used as part of efficiency ratios like cycles per instruction (CPI). Processor clocks may stop ticking under circumstances like the following:

- The processor is halted when there is nothing for the CPU to do. For example, the processor may halt to save power while the computer is servicing an I/O request. When Intel Hyper-Threading Technology is enabled, both logical processors must be halted for performance-monitoring counters to be powered down.
- The processor is asleep as a result of being halted or because of a power-management scheme. There are different levels of sleep. In the some deep sleep levels, the time-stamp counter stops counting.

In addition, processor core clocks may undergo transitions at different ratios relative to the processor’s bus clock frequency. Some of the situations that can cause processor core clock to undergo frequency transitions include:

- TM2 transitions.
- Enhanced Intel SpeedStep Technology transitions (P-state transitions).

For Intel processors that support TM2, the processor core clocks may operate at a frequency that differs from the Processor Base frequency (as indicated by processor frequency information reported by CPUID instruction). See Section 18.7.2 for more detail.

Due to the above considerations there are several important clocks referenced in this manual:

- **Base Clock** — The frequency of this clock is the frequency of the processor when the processor is not in turbo mode, and not being throttled via Intel SpeedStep.
- **Maximum Clock** — This is the maximum frequency of the processor when turbo mode is at the highest point.
- **Bus Clock** — These clockticks increment at a fixed frequency and help coordinate the bus on some systems.

- **Core Crystal Clock** — This is a clock that runs at fixed frequency; it coordinates the clocks on all packages across the system.
- **Non-halted Clockticks** — Measures clock cycles in which the specified logical processor is not halted and is not in any power-saving state. When Intel Hyper-Threading Technology is enabled, ticks can be measured on a per-logical-processor basis. There are also performance events on dual-core processors that measure clockticks per logical processor when the processor is not halted.
- **Non-sleep Clockticks** — Measures clock cycles in which the specified physical processor is not in a sleep mode or in a power-saving state. These ticks cannot be measured on a logical-processor basis.
- **Time-stamp Counter** — See Section 17.17, “Time-Stamp Counter”.
- **Reference Clockticks** — TM2 or Enhanced Intel SpeedStep technology are two examples of processor features that can cause processor core clockticks to represent non-uniform tick intervals due to change of bus ratios. Performance events that counts clockticks of a constant reference frequency was introduced Intel Core Duo and Intel Core Solo processors. The mechanism is further enhanced on processors based on Intel Core microarchitecture.

Some processor models permit clock cycles to be measured when the physical processor is not in deep sleep (by using the time-stamp counter and the RDTSC instruction). Note that such ticks cannot be measured on a per-logical-processor basis. See Section 17.17, “Time-Stamp Counter,” for detail on processor capabilities.

The first two methods use performance counters and can be set up to cause an interrupt upon overflow (for sampling). They may also be useful where it is easier for a tool to read a performance counter than to use a time stamp counter (the timestamp counter is accessed using the RDTSC instruction).

For applications with a significant amount of I/O, there are two ratios of interest:

- **Non-halted CPI** — Non-halted clockticks/instructions retired measures the CPI for phases where the CPU was being used. This ratio can be measured on a logical-processor basis when Intel Hyper-Threading Technology is enabled.
- **Nominal CPI** — Time-stamp counter ticks/instructions retired measures the CPI over the duration of a program, including those periods when the machine halts while waiting for I/O.

### 18.7.1 Non-Halted Reference Clockticks

Software can use UnHalted Reference Cycles on either a general purpose performance counter using event mask 0x3C and umask 0x01 or on fixed function performance counter 2 to count at a constant rate. These events count at a consistent rate irrespective of P-state, TM2, or frequency transitions that may occur to the processor. The UnHalted Reference Cycles event may count differently on the general purpose event and fixed counter.

### 18.7.2 Cycle Counting and Opportunistic Processor Operation

As a result of the state transitions due to opportunistic processor performance operation (see Chapter 14, “Power and Thermal Management”), a logical processor or a processor core can operate at frequency different from the Processor Base frequency.

The following items are expected to hold true irrespective of when opportunistic processor operation causes state transitions:

- The time stamp counter operates at a fixed-rate frequency of the processor.
- The IA32\_MPERF counter increments at a fixed frequency irrespective of any transitions caused by opportunistic processor operation.
- The IA32\_FIXED\_CTR2 counter increments at the same TSC frequency irrespective of any transitions caused by opportunistic processor operation.
- The Local APIC timer operation is unaffected by opportunistic processor operation.
- The TSC, IA32\_MPERF, and IA32\_FIXED\_CTR2 operate at close to the maximum non-turbo frequency, which is equal to the product of scalable bus frequency and maximum non-turbo ratio.

### 18.7.3 Determining the Processor Base Frequency

For Intel processors in which the nominal core crystal clock frequency is enumerated in CPUID.15H.ECX and the core crystal clock ratio is encoded in CPUID.15H (see Table 3-8 “Information Returned by CPUID Instruction”), the nominal TSC frequency can be determined by using the following equation:

$$\text{Nominal TSC frequency} = (\text{CPUID.15H.ECX}[31:0] * \text{CPUID.15H.EBX}[31:0]) \div \text{CPUID.15H.EAX}[31:0]$$

For Intel processors in which CPUID.15H.EBX[31:0] ÷ CPUID.0x15.EAX[31:0] is enumerated but CPUID.15H.ECX is not enumerated, Table 18-75 can be used to look up the nominal core crystal clock frequency.

**Table 18-75. Nominal Core Crystal Clock Frequency**

| Processor Families/Processor Number Series <sup>1</sup>   | Nominal Core Crystal Clock Frequency |
|---|--------------------------------------|
| Future Intel® Xeon® processors with CPUID signature 06_55H.   | 25 MHz                               |
| 6th and 7th generation Intel® Core™ processors (does not include Intel® Xeon® processors).  | 24 MHz                               |
| Next Generation Intel® Atom™ processors based on Goldmont Microarchitecture with CPUID signature 06_5CH (does not include Intel Xeon processors). | 19.2 MHz                             |

**NOTES:**

1. For any processor in which CPUID.15H is enumerated and MSR\_PLATFORM\_INFO[15:8] (which gives the scalable bus frequency) is available, a more accurate frequency can be obtained by using CPUID.15H.

#### 18.7.3.1 For Intel® Processors Based on Microarchitecture Code Name Sandy Bridge, Ivy Bridge, Haswell and Broadwell

The scalable bus frequency is encoded in the bit field MSR\_PLATFORM\_INFO[15:8] and the nominal TSC frequency can be determined by multiplying this number by a bus speed of 100 MHz.

#### 18.7.3.2 For Intel® Processors Based on Microarchitecture Code Name Nehalem

The scalable bus frequency is encoded in the bit field MSR\_PLATFORM\_INFO[15:8] and the nominal TSC frequency can be determined by multiplying this number by a bus speed of 133.33 MHz.

#### 18.7.3.3 For Intel® Atom™ Processors Based on the Silvermont Microarchitecture (Including Intel Processors Based on Airmont Microarchitecture)

The scalable bus frequency is encoded in the bit field MSR\_PLATFORM\_INFO[15:8] and the nominal TSC frequency can be determined by multiplying this number by the scalable bus frequency. The scalable bus frequency is encoded in the bit field MSR\_FSB\_FREQ[2:0] for Intel Atom processors based on the Silvermont microarchitecture, and in bit field MSR\_FSB\_FREQ[3:0] for processors based on the Airmont microarchitecture; see Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*.



### 18.7.3.4 For Intel® Core™ 2 Processor Family and for Intel® Xeon® Processors Based on Intel Core Microarchitecture

For processors based on Intel Core microarchitecture, the scalable bus frequency is encoded in the bit field MSR\_FSB\_FREQ[2:0] at (0CDH), see Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*. The maximum resolved bus ratio can be read from the following bit field:

- If XE operation is disabled, the maximum resolved bus ratio can be read in MSR\_PLATFORM\_ID[12:8]. It corresponds to the Processor Base frequency.
- IF XE operation is enabled, the maximum resolved bus ratio is given in MSR\_PERF\_STATUS[44:40], it corresponds to the maximum XE operation frequency configured by BIOS.

XE operation of an Intel 64 processor is implementation specific. XE operation can be enabled only by BIOS. If MSR\_PERF\_STATUS[31] is set, XE operation is enabled. The MSR\_PERF\_STATUS[31] field is read-only.

## 18.8 IA32\_PERF\_CAPABILITIES MSR ENUMERATION

The layout of IA32\_PERF\_CAPABILITIES MSR is shown in Figure 18-63, it provides enumeration of a variety of interfaces:

- IA32\_PERF\_CAPABILITIES.LBR\_FMT[bits 5:0]: encodes the LBR format, details are described in Section 17.4.8.1.
- IA32\_PERF\_CAPABILITIES.PEBSTrap[6]: Trap/Fault-like indicator of PEBS recording assist, see Section 18.6.2.4.2.
- IA32\_PERF\_CAPABILITIES.PEBSArchRegs[7]: Indicator of PEBS assist save architectural registers, see Section 18.6.2.4.2.
- IA32\_PERF\_CAPABILITIES.PEBS\_FMT[bits 11:8]: Specifies the encoding of the layout of PEBS records, see Section 18.6.2.4.2.
- IA32\_PERF\_CAPABILITIES.FREEZE\_WHILE\_SMM[12]: Indicates IA32\_DEBUGCTL.FREEZE\_WHILE\_SMM is supported if 1, see Section 18.8.1.
- IA32\_PERF\_CAPABILITIES.FULL\_WRITE[13]: Indicates the processor supports IA32\_A\_PMCx interface for updating bits 32 and above of IA32\_PMCx, see Section 18.2.5.




Figure 18-63. Layout of IA32\_PERF\_CAPABILITIES MSR

### 18.8.1 Filtering of SMM Handler Overhead

When performance monitoring facilities and/or branch profiling facilities (see Section 17.5, “Last Branch, Interrupt, and Exception Recording (Intel® Core™ 2 Duo and Intel® Atom™ Processors)”) are enabled, these facilities capture event counts, branch records and branch trace messages occurring in a logical processor. The occurrence of interrupts, instruction streams due to various interrupt handlers all contribute to the results recorded by these facilities.

If CPUID.01H:ECX.PDCM[bit 15] is 1, the processor supports the IA32\_PERF\_CAPABILITIES MSR. If IA32\_PERF\_CAPABILITIES.FREEZE\_WHILE\_SMM[Bit 12] is 1, the processor supports the ability for system software using performance monitoring and/or branch profiling facilities to filter out the effects of servicing system management interrupts.

If the FREEZE\_WHILE\_SMM capability is enabled on a logical processor and after an SMI is delivered, the processor will clear all the enable bits of IA32\_PERF\_GLOBAL\_CTRL, save a copy of the content of IA32\_DEBUGCTL and disable LBR, BTF, TR, and BTS fields of IA32\_DEBUGCTL before transferring control to the SMI handler.

The enable bits of IA32\_PERF\_GLOBAL\_CTRL will be set to 1, the saved copy of IA32\_DEBUGCTL prior to SMI delivery will be restored, after the SMI handler issues RSM to complete its servicing.

It is the responsibility of the SMM code to ensure the state of the performance monitoring and branch profiling facilities are preserved upon entry or until prior to exiting the SMM. If any of this state is modified due to actions by the SMM code, the SMM code is required to restore such state to the values present at entry to the SMM handler.

System software is allowed to set IA32\_DEBUGCTL.FREEZE\_WHILE\_SMM[bit 14] to 1 only supported as indicated by IA32\_PERF\_CAPABILITIES.FREEZE\_WHILE\_SMM[Bit 12] reporting 1.

## 12. Updates to Chapter 19, Volume 3B

Change bars show changes to Chapter 19 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2*.

-----  
Changes to this chapter: Added note on availability of Performance Monitoring Events Document and location. Updates to add support for 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series to existing sections.

### NOTE

The event tables listed in this chapter provide information for tool developers to support architectural and model-specific performance monitoring events. The tables are up to date at processor launch, but are subject to changes. The most up to date event tables and additional details of performance event implementation can be found here:

1) In the document titled "*Intel® 64 and IA32 Architectures Performance Monitoring Events*", Document number: 335279, located here: [https://software.intel.com/sites/default/files/managed/8b/6e/335279\\_performance\\_monitoring\\_events\\_guide.pdf](https://software.intel.com/sites/default/files/managed/8b/6e/335279_performance_monitoring_events_guide.pdf).

These event tables are also available on the web and are located at: <https://download.01.org/perfmon/index/>.

2) Performance monitoring event files for Intel processors are hosted by the Intel Open Source Technology Center. These files can be downloaded here: <https://download.01.org/perfmon/>.

This chapter lists the performance monitoring events that can be monitored with the Intel 64 or IA-32 processors. The ability to monitor performance events and the events that can be monitored in these processors are mostly model-specific, except for architectural performance events, described in Section 19.1.

Model-specific performance events are listed for each generation of microarchitecture:

- Section 19.2 - Processors based on Skylake microarchitecture
- Section 19.3 - Processors based on Skylake, Kaby Lake and Coffee Lake microarchitectures
- Section 19.4 - Processors based on Knights Landing and Knights Mill microarchitectures
- Section 19.5 - Processors based on Broadwell microarchitecture
- Section 19.6 - Processors based on Haswell microarchitecture
- Section 19.6.1 - Processors based on Haswell-E microarchitecture
- Section 19.7 - Processors based on Ivy Bridge microarchitecture
- Section 19.7.1 - Processors based on Ivy Bridge-E microarchitecture
- Section 19.8 - Processors based on Sandy Bridge microarchitecture
- Section 19.9 - Processors based on Intel® microarchitecture code name Nehalem
- Section 19.10 - Processors based on Intel® microarchitecture code name Westmere
- Section 19.11 - Processors based on Enhanced Intel® Core™ microarchitecture
- Section 19.12 - Processors based on Intel® Core™ microarchitecture
- Section 19.13 - Processors based on the Goldmont microarchitecture
- Section 19.15 - Processors based on the Silvermont microarchitecture
- Section 19.15.1 - Processors based on the Airmont microarchitecture
- Section 19.16 - 45 nm and 32 nm Intel® Atom™ Processors
- Section 19.17 - Intel® Core™ Solo and Intel® Core™ Duo processors
- Section 19.18 - Processors based on Intel NetBurst® microarchitecture
- Section 19.19 - Pentium® M family processors
- Section 19.20 - P6 family processors
- Section 19.21 - Pentium® processors

**NOTE**

These performance monitoring events are intended to be used as guides for performance tuning. The counter values reported by the performance monitoring events are approximate and believed to be useful as relative guides for tuning software. Known discrepancies are documented where applicable.

All performance event encodings not documented in the appropriate tables for the given processor are considered reserved, and their use will result in undefined counter updates with associated overflow actions.

## 19.1 ARCHITECTURAL PERFORMANCE MONITORING EVENTS

Architectural performance events are introduced in Intel Core Solo and Intel Core Duo processors. They are also supported on processors based on Intel Core microarchitecture. Table 19-1 lists pre-defined architectural performance events that can be configured using general-purpose performance counters and associated event-select registers.

**Table 19-1. Architectural Performance Events**

| Event Num. | Event Mask Name                        | Umask Value | Description   |
|------------|--|-------------|---|
| 3CH        | UnHalted Core Cycles                   | 00H         | Counts core clock cycles whenever the logical processor is in C0 state (not halted). The frequency of this event varies with state transitions in the core. |
| 3CH        | UnHalted Reference Cycles <sup>1</sup> | 01H         | Counts at a fixed frequency whenever the logical processor is in C0 state (not halted).   |
| C0H        | Instructions Retired                   | 00H         | Counts when the last uop of an instruction retires.   |
| 2EH        | LLC Reference                          | 4FH         | Counts requests originating from the core that reference a cache line in the last level on-die cache.   |
| 2EH        | LLC Misses                             | 41H         | Counts each cache miss condition for references to the last level on-die cache.   |
| C4H        | Branch Instruction Retired             | 00H         | Counts when the last uop of a branch instruction retires.   |
| C5H        | Branch Misses Retired                  | 00H         | Counts when the last uop of a branch instruction retires which corrected misprediction of the branch prediction hardware at execution time.                 |

**NOTES:**

1. Current implementations count at core crystal clock, TSC, or bus clock frequency.

Fixed-function performance counters count only events defined in Table 19-2.

**Table 19-2. Fixed-Function Performance Counter and Pre-defined Performance Events**

| Fixed-Function Performance Counter | Address | Event Mask Mnemonic | Description   |
|------------------------------------|---------|---------------------|---|
| IA32_PERF_FIXED_CTR0               | 309H    | Inst_Retired.Any    | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers. |

**Table 19-2. Fixed-Function Performance Counter and Pre-defined Performance Events (Contd.)**

| Fixed-Function Performance Counter | Address | Event Mask Mnemonic   | Description   |
|------------------------------------|---------|---|---|
| IA32_PERF_FIXED_CTR1               | 30AH    | CPU_CLK_UNHALTED.THREAD/CPU_CLK_UNHALTED.CORE/CPU_CLK_UNHALTED.THREAD_ANY | <p>The CPU_CLK_UNHALTED.THREAD event counts the number of core cycles while the logical processor is not in a halt state.</p> <p>If there is only one logical processor in a processor core, CPU_CLK_UNHALTED.CORE counts the unhalted cycles of the processor core.</p> <p>If there are more than one logical processor in a processor core, CPU_CLK_UNHALTED.THREAD_ANY is supported by programming IA32_FIXED_CTR_CTRL[bit 6]AnyThread = 1.</p> <p>The core frequency may change from time to time due to transitions associated with Enhanced Intel SpeedStep Technology or TM2. For this reason this event may have a changing ratio with regards to time.</p> |
| IA32_PERF_FIXED_CTR2               | 30BH    | CPU_CLK_UNHALTED.REF_TSC  | <p>This event counts the number of reference cycles at the TSC rate when the core is not in a halt state and not in a TM stop-clock state. The core enters the halt state when it is running the HLT instruction or the MWAIT instruction. This event is not affected by core frequency changes (e.g., P states) but counts at the same frequency as the time stamp counter. This event can approximate elapsed time while the core was not in a halt state and not in a TM stopclock state.</p>  |

## 19.2 PERFORMANCE MONITORING EVENTS FOR INTEL® XEON® PROCESSOR SCALABLE FAMILY

The Intel® Xeon® Processor Scalable Family is based on the Skylake microarchitecture. These processors support the architectural performance monitoring events listed in Table 19-1. Fixed counters in the core PMU support the architecture events defined in Table 19-2. Model-specific performance monitoring events in the processor core are listed in Table 19-4. The events in Table 19-4 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following value: 06\_55H .

The comment column in Table 19-4 uses abbreviated letters to indicate additional conditions applicable to the Event Mask Mnemonic. For event umasks listed in Table 19-4 that do not show "AnyT", users should refrain from programming "AnyThread =1" in IA32\_PERF\_EVTSELx.

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic           | Description  | Comment       |
|------------|-------------|-------------------------------|--|---------------|
| 00H        | 01H         | INST_RETIREDA.ANY             | Counts the number of instructions retired from execution. For instructions that consist of multiple micro-ops, Counts the retirement of the last micro-op of the instruction. Counting continues during hardware interrupts, traps, and inside interrupt handlers. Notes: INST_RETIREDA.ANY is counted by a designated fixed counter, leaving the four (eight when Hyperthreading is disabled) programmable counters available for other events. INST_RETIREDA.ANY_P is counted by a programmable counter and it is an architectural performance event. Counting: Faulting executions of GETSEC/VM entry/VM Exit/MWait will not count as retired instructions.                                 | Fixed Counter |
| 00H        | 02H         | CPU_CLK_UNHALTED.THREAD       | Counts the number of core cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. This event is a component in many key event ratios. The core frequency may change from time to time due to transitions associated with Enhanced Intel SpeedStep Technology or TM2. For this reason this event may have a changing ratio with regards to time. When the core frequency is constant, this event can approximate elapsed time while the core was not in the halt state. It is counted on a dedicated fixed counter, leaving the four (eight when Hyperthreading is disabled) programmable counters available for other events. | Fixed Counter |
| 00H        | 02H         | CPU_CLK_UNHALTED.THREAD_A.ANY | Core cycles when at least one thread on the physical core is not in halt state.  | AnyThread=1   |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description  | Comment       |
|------------|-------------|---------------------------------|--|---------------|
| 00H        | 03H         | CPU_CLK_UNHALTED.REF_TSC        | Counts the number of reference cycles when the core is not in a halt state. The core enters the halt state when it is running the HLT instruction or the MWAIT instruction. This event is not affected by core frequency changes (for example, P states, TM2 transitions) but has the same incrementing frequency as the time stamp counter. This event can approximate elapsed time while the core was not in a halt state. This event has a constant ratio with the CPU_CLK_UNHALTED.REF_XCLK event. It is counted on a dedicated fixed counter, leaving the four (eight when Hyperthreading is disabled) programmable counters available for other events. Note: On all current platforms this event stops counting during 'throttling (TM)' states duty off periods the processor is 'halted'. The counter update is done at a lower clock rate than the core clock the overflow status bit for this counter may appear 'sticky'. After the counter has overflowed and software clears the overflow status bit and resets the counter to less than MAX. The reset value to the counter is not clocked immediately so the overflow status bit will flip "high (1)" and generate another PMI (if enabled) after which the reset value gets clocked into the counter. Therefore, software will get the interrupt, read the overflow status bit '1' for bit 34 while the counter value is less than MAX. Software should ignore this case. | Fixed Counter |
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD         | Counts how many times the load operation got the true Block-on-Store blocking code preventing store forwarding. This includes cases when: a. preceding store conflicts with the load (incomplete overlap), b. store forwarding is impossible due to u-arch limitations, c. preceding lock RMW operations are not forwarded, d. store has the no-forward bit set (uncacheable/page-split/masked stores), e. all-blocking stores are used (mostly, fences and port I/O), and others. The most common case is a load blocked due to its address range overlapping with a preceding smaller uncompleted store. Note: This event does not take into account cases of out-of-SW-control (for example, SbTailHit), unknown physical STA, and cases of blocking loads on store due to being non-WB memory type or a lock. These cases are covered by other events. See the table of not supported store forwards in the Optimization Guide.  |               |
| 03H        | 08H         | LD_BLOCKS.NO_SR                 | The number of times that split load operations are temporarily blocked because all resources for handling the split accesses are in use.   |               |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS | Counts false dependencies in MOB when the partial comparison upon loose net check and dependency was resolved by the Enhanced Loose net mechanism. This may not result in high performance penalties. Loose net checks can fail when loads and stores are 4k aliased.  |               |



**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                   | Description   | Comment                              |
|------------|-------------|---------------------------------------|---|--------------------------------------|
| 08H        | 01H         | DTLB_LOAD_MISSES.MISS_CAUSES_A_WALK   | Counts demand data loads that caused a page walk of any page size (4K/2M/4M/1G). This implies it missed in all TLB levels, but the walk need not have completed.  |                                      |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED_4K    | Counts demand data loads that caused a completed page walk (4K page size). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |                                      |
| 08H        | 04H         | DTLB_LOAD_MISSES.WALK_COMPLETED_2M_4M | Counts demand data loads that caused a completed page walk (2M and 4M page sizes). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |                                      |
| 08H        | 08H         | DTLB_LOAD_MISSES.WALK_COMPLETED_1G    | Counts load misses in all DTLB levels that cause a completed page walk (1G page size). The page walk can end with or without a fault.   |                                      |
| 08H        | 0EH         | DTLB_LOAD_MISSES.WALK_COMPLETED       | Counts demand data loads that caused a completed page walk of any page size (4K/2M/4M/1G). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |                                      |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_PENDING         | Counts 1 per cycle for each PMH that is busy with a page walk for a load. EPT page walk duration are excluded in Skylake microarchitecture.   |                                      |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_ACTIVE          | Counts cycles when at least one PMH (Page Miss Handler) is busy with a page walk for a load.  | CounterMask=1<br>CMSK1               |
| 08H        | 20H         | DTLB_LOAD_MISSES.STLB_HIT             | Counts loads that miss the DTLB (Data TLB) and hit the STLB (Second level TLB).   |                                      |
| 0DH        | 01H         | INT_MISC.RECOVERY_CYCLES              | Core cycles the Resource allocator was stalled due to recovery from an earlier branch misprediction or machine clear event.   |                                      |
| 0DH        | 01H         | INT_MISC.RECOVERY_CYCLES_ANY          | Core cycles the allocator was stalled due to recovery from earlier clear event for any thread running on the physical core (e.g. misprediction or memory nuke).   | AnyThread=1 AnyT                     |
| 0DH        | 80H         | INT_MISC.CLEAR_RESTEER_CYCLES         | Cycles the issue-stage is waiting for front-end to fetch from resteeered path following branch misprediction or machine clear events.   |                                      |
| 0EH        | 01H         | UOPS_ISSUED.ANY                       | Counts the number of uops that the Resource Allocation Table (RAT) issues to the Reservation Station (RS).  |                                      |
| 0EH        | 01H         | UOPS_ISSUED.STALL_CYCLES              | Counts cycles during which the Resource Allocation Table (RAT) does not issue any uops to the reservation station (RS) for the current thread.  | CounterMask=1<br>Invert=1 CMSK1, INV |
| 0EH        | 02H         | UOPS_ISSUED.VECTOR_WIDTH_MISMATCH     | Counts the number of Blend Uops issued by the Resource Allocation Table (RAT) to the reservation station (RS) in order to preserve upper bits of vector registers. Starting with the Skylake microarchitecture, these Blend uops are needed since every Intel SSE instruction executed in Dirty Upper State needs to preserve bits 128-255 of the destination register. For more information, refer to Mixing Intel AVX and Intel SSE Code section of the Optimization Guide. |                                      |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment       |
|------------|-------------|--------------------------------|--|---------------|
| 0EH        | 20H         | UOPS_ISSUED.SLOW_LEA           | Number of slow LEA uops being allocated. A uop is generally considered SlowLea if it has 3 sources (e.g. 2 sources + immediate) regardless if as a result of LEA instruction or not. |               |
| 14H        | 01H         | ARITH.DIVIDER_ACTIVE           | Cycles when divide unit is busy executing divide or square root operations. Accounts for integer and floating-point operations.  | CounterMask=1 |
| 24H        | 21H         | L2_RQSTS.DEMAND_DATA_RD_MISS   | Counts the number of demand Data Read requests that miss L2 cache. Only not rejected loads are counted.  |               |
| 24H        | 22H         | L2_RQSTS.RFO_MISS              | Counts the RFO (Read-for-Ownership) requests that miss L2 cache.   |               |
| 24H        | 24H         | L2_RQSTS.CODE_RD_MISS          | Counts L2 cache misses when fetching instructions.   |               |
| 24H        | 27H         | L2_RQSTS.ALL_DEMAND_MISS       | Demand requests that miss L2 cache.  |               |
| 24H        | 38H         | L2_RQSTS.PF_MISS               | Counts requests from the L1/L2/L3 hardware prefetchers or Load software prefetches that miss L2 cache.   |               |
| 24H        | 3FH         | L2_RQSTS.MISS                  | All requests that miss L2 cache.   |               |
| 24H        | 41H         | L2_RQSTS.DEMAND_DATA_RD_HIT    | Counts the number of demand Data Read requests that hit L2 cache. Only non rejected loads are counted.   |               |
| 24H        | 42H         | L2_RQSTS.RFO_HIT               | Counts the RFO (Read-for-Ownership) requests that hit L2 cache.  |               |
| 24H        | 44H         | L2_RQSTS.CODE_RD_HIT           | Counts L2 cache hits when fetching instructions, code reads.   |               |
| 24H        | D8H         | L2_RQSTS.PF_HIT                | Counts requests from the L1/L2/L3 hardware prefetchers or Load software prefetches that hit L2 cache.  |               |
| 24H        | E1H         | L2_RQSTS.ALL_DEMAND_DATA_RD    | Counts the number of demand Data Read requests (including requests from L1D hardware prefetchers). These loads may hit or miss L2 cache. Only non rejected loads are counted.        |               |
| 24H        | E2H         | L2_RQSTS.ALL_RFO               | Counts the total number of RFO (read for ownership) requests to L2 cache. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches.                          |               |
| 24H        | E4H         | L2_RQSTS.ALL_CODE_RD           | Counts the total number of L2 code requests.   |               |
| 24H        | E7H         | L2_RQSTS.ALL_DEMAND_REFERENCES | Demand requests to L2 cache.   |               |
| 24H        | F8H         | L2_RQSTS.ALL_PF                | Counts the total number of requests from the L2 hardware prefetchers.  |               |
| 24H        | FFH         | L2_RQSTS.REFERENCES            | All L2 requests.   |               |
| 28H        | 07H         | CORE_POWER.LVLO_TURBO_LICENSE  | Core cycles where the core was running with power-delivery for baseline license level 0. This includes non-AVX codes, SSE, AVX 128-bit, and low-current AVX 256-bit codes.           |               |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                           | Description   | Comment                       |
|------------|-------------|---|---|-------------------------------|
| 28H        | 18H         | CORE_POWER.LVL1_TURBO_LIC<br>ENSE             | Core cycles where the core was running with power-delivery for license level 1. This includes high current AVX 256-bit instructions as well as low current AVX 512-bit instructions.  |                               |
| 28H        | 20H         | CORE_POWER.LVL2_TURBO_LIC<br>ENSE             | Core cycles where the core was running with power-delivery for license level 2 (introduced in Skylake Server microarchitecture). This includes high current AVX 512-bit instructions.   |                               |
| 28H        | 40H         | CORE_POWER.THROTTLE                           | Core cycles the out-of-order engine was throttled due to a pending power level request.   |                               |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS                        | Counts core-originated cacheable requests that miss the L3 cache (Longest Latency cache). Requests include data and code reads, Reads-for-Ownership (RFOs), speculative accesses and hardware prefetches from L1 and L2. It does not include all misses to the L3.  | See Table 19-1.               |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFEREN<br>CE               | Counts core-originated cacheable requests to the L3 cache (Longest Latency cache). Requests include data and code reads, Reads-for-Ownership (RFOs), speculative accesses and hardware prefetches from L1 and L2. It does not include all accesses to the L3.   | See Table 19-1.               |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_<br>P                 | This is an architectural event that counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. For this reason, this event may have a changing ratio with regards to wall clock time. | See Table 19-1.               |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_<br>P_ANY             | Core cycles when at least one thread on the physical core is not in halt state.   | AnyThread=1 AnyT              |
| 3CH        | 00H         | CPU_CLK_UNHALTED.RINGO_TR<br>ANS              | Counts when the Current Privilege Level (CPL) transitions from ring 1, 2 or 3 to ring 0 (Kernel).   | EdgeDetect=1<br>CounterMask=1 |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.<br>REF_XCLK          | Core crystal clock cycles when the thread is unhalting.   | See Table 19-1.               |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.<br>REF_XCLK_ANY      | Core crystal clock cycles when at least one thread on the physical core is unhalting.   | AnyThread=1 AnyT              |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF_XCLK                     | Core crystal clock cycles when the thread is unhalting.   | See Table 19-1.               |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF_XCLK<br>_ANY             | Core crystal clock cycles when at least one thread on the physical core is unhalting.   | AnyThread=1 AnyT              |
| 3CH        | 02H         | CPU_CLK_THREAD_UNHALTED.<br>ONE_THREAD_ACTIVE | Core crystal clock cycles when this thread is unhalting and the other thread is halted.   |                               |
| 3CH        | 02H         | CPU_CLK_UNHALTED.ONE_THR<br>EAD_ACTIVE        | Core crystal clock cycles when this thread is unhalting and the other thread is halted.   |                               |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description  | Comment                                     |
|------------|-------------|--|--|---|
| 48H        | 01H         | L1D_PEND_MISS.PENDING                  | Counts duration of L1D miss outstanding, that is each cycle number of Fill Buffers (FB) outstanding required by Demand Reads. FB either is held by demand loads, or it is held by non-demand loads and gets hit at least once by demand. The valid outstanding interval is defined until the FB deallocation by one of the following ways: from FB allocation, if FB is allocated by demand from the demand Hit FB, if it is allocated by hardware or software prefetch. Note: In the L1D, a Demand Read contains cacheable or noncacheable demand loads, including ones causing cache-line splits and reads due to page walks resulted from any request type. |   |
| 48H        | 01H         | L1D_PEND_MISS.PENDING_CYCLES           | Counts duration of L1D miss outstanding in cycles.   | CounterMask=1<br>CMSK1                      |
| 48H        | 01H         | L1D_PEND_MISS.PENDING_CYCLES_ANY       | Cycles with L1D load Misses outstanding from any thread on physical core.  | CounterMask=1<br>AnyThread=1<br>CMSK1, AnyT |
| 48H        | 02H         | L1D_PEND_MISS.FB_FULL                  | Number of times a request needed a FB (Fill Buffer) entry but there was no entry available for it. A request includes cacheable/uncacheable demands that are load, store or SW prefetch instructions.  |   |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK   | Counts demand data stores that caused a page walk of any page size (4K/2M/4M/1G). This implies it missed in all TLB levels, but the walk need not have completed.  |   |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED_4K    | Counts demand data stores that caused a completed page walk (4K page size). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |   |
| 49H        | 04H         | DTLB_STORE_MISSES.WALK_COMPLETED_2M_4M | Counts demand data stores that caused a completed page walk (2M and 4M page sizes). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |   |
| 49H        | 08H         | DTLB_STORE_MISSES.WALK_COMPLETED_1G    | Counts store misses in all DTLB levels that cause a completed page walk (1G page size). The page walk can end with or without a fault.   |   |
| 49H        | 0EH         | DTLB_STORE_MISSES.WALK_COMPLETED       | Counts demand data stores that caused a completed page walk of any page size (4K/2M/4M/1G). This implies it missed in all TLB levels. The page walk can end with or without a fault.   |   |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_PENDING         | Counts 1 per cycle for each PMH that is busy with a page walk for a store. EPT page walk duration are excluded in Skylake microarchitecture.   |   |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_ACTIVE          | Counts cycles when at least one PMH (Page Miss Handler) is busy with a page walk for a store.  | CounterMask=1<br>CMSK1                      |
| 49H        | 20H         | DTLB_STORE_MISSES.STLB_HIT             | Stores that miss the DTLB (Data TLB) and hit the STLB (2nd Level TLB).   |   |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description  | Comment  |
|------------|-------------|---|--|--|
| 4CH        | 01H         | LOAD_HIT_PRE.SW_PF  | Counts all not software-prefetch load dispatches that hit the fill buffer (FB) allocated for the software prefetch. It can also be incremented by some lock instructions. So it should only be used with profiling so that the locks can be excluded by ASM (Assembly File) inspection of the nearby instructions. |  |
| 4FH        | 10H         | EPT.WALK_PENDING  | Counts cycles for each PMH (Page Miss Handler) that is busy with an EPT (Extended Page Table) walk for any request type.   |  |
| 51H        | 01H         | L1D.REPLACEMENT   | Counts L1D data line replacements including opportunistic replacements, and replacements that require stall-for-replace or block-for-replace.  |  |
| 54H        | 01H         | TX_MEM.ABORT_CONFLICT   | Number of times a TSX line had a cache conflict.   |  |
| 54H        | 02H         | TX_MEM.ABORT_CAPACITY   | Number of times a transactional abort was signaled due to a data capacity limitation for transactional reads or writes.  |  |
| 54H        | 04H         | TX_MEM.ABORT_HLE_STORE_T<br>O_ELIDED_LOCK                     | Number of times a TSX Abort was triggered due to a non-release/commit store to lock.   |  |
| 54H        | 08H         | TX_MEM.ABORT_HLE_ELISION_<br>BUFFER_NOT_EMPTY                 | Number of times a TSX Abort was triggered due to commit but Lock Buffer not empty.   |  |
| 54H        | 10H         | TX_MEM.ABORT_HLE_ELISION_<br>BUFFER_MISMATCH                  | Number of times a TSX Abort was triggered due to release/commit but data and address mismatch.   |  |
| 54H        | 20H         | TX_MEM.ABORT_HLE_ELISION_<br>BUFFER_UNSUPPORTED_ALIGN<br>MENT | Number of times a TSX Abort was triggered due to attempting an unsupported alignment from Lock Buffer.   |  |
| 54H        | 40H         | TX_MEM.HLE_ELISION_BUFFER_<br>FULL                            | Number of times we could not allocate Lock Buffer.   |  |
| 5DH        | 01H         | TX_EXEC.MISC1   | Unfriendly TSX abort triggered by a flowmarker.  |  |
| 5DH        | 02H         | TX_EXEC.MISC2   | Unfriendly TSX abort triggered by a vzeroupper instruction.  |  |
| 5DH        | 04H         | TX_EXEC.MISC3   | Unfriendly TSX abort triggered by a nest count that is too deep.   |  |
| 5DH        | 08H         | TX_EXEC.MISC4   | RTM region detected inside HLE.  |  |
| 5DH        | 10H         | TX_EXEC.MISC5   | Counts the number of times an HLE XACQUIRE instruction was executed inside an RTM transactional region.  |  |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES  | Counts cycles during which the reservation station (RS) is empty for the thread.; Note: In ST-mode, not active thread should drive 0. This is usually caused by severely costly branch mispredictions, or allocator/FE issues.   |  |
| 5EH        | 01H         | RS_EVENTS.EMPTY_END   | Counts end of periods where the Reservation Station (RS) was empty. Could be useful to precisely locate front-end Latency Bound issues.  | EdgeDetect=1<br>CounterMask=1<br>Invert=1 CMSK1, INV |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                                     | Description   | Comment                |
|------------|-------------|---|---|------------------------|
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD             | Counts the number of offcore outstanding Demand Data Read transactions in the super queue (SQ) every cycle. A transaction is considered to be in the Offcore outstanding state between L2 miss and transaction completion sent to requestor. See the corresponding Umask under OFFCORE_REQUESTS. Note: A prefetch promoted to Demand is counted from the promotion point. |                        |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_DATA_RD | Counts cycles when offcore outstanding Demand Data Read transactions are present in the super queue (SQ). A transaction is considered to be in the Offcore outstanding state between L2 miss and transaction completion sent to requestor (SQ de-allocation).   | CounterMask=1<br>CMSK1 |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD_GE_6        | Cycles with at least 6 offcore outstanding Demand Data Read transactions in uncore queue.   | CounterMask=6<br>CMSK6 |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_CODE_RD             | Counts the number of offcore outstanding Code Reads transactions in the super queue every cycle. The 'Offcore outstanding' state of the transaction lasts from the L2 miss until the sending transaction completion to requestor (SQ deallocation). See the corresponding Umask under OFFCORE_REQUESTS.   | CounterMask=1<br>CMSK1 |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_CODE_RD | Counts the number of offcore outstanding Code Reads transactions in the super queue every cycle. The 'Offcore outstanding' state of the transaction lasts from the L2 miss until the sending transaction completion to requestor (SQ deallocation). See the corresponding Umask under OFFCORE_REQUESTS.   | CMSK1                  |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO                 | Counts the number of offcore outstanding RFO (store) transactions in the super queue (SQ) every cycle. A transaction is considered to be in the Offcore outstanding state between L2 miss and transaction completion sent to requestor (SQ de-allocation). See corresponding Umask under OFFCORE_REQUESTS.  | CounterMask=1<br>CMSK1 |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_RFO     | Counts the number of offcore outstanding demand rfo Reads transactions in the super queue every cycle. The 'Offcore outstanding' state of the transaction lasts from the L2 miss until the sending transaction completion to requestor (SQ deallocation). See the corresponding Umask under OFFCORE_REQUESTS.   | CMSK1                  |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD                | Counts the number of offcore outstanding cacheable Core Data Read transactions in the super queue every cycle. A transaction is considered to be in the Offcore outstanding state between L2 miss and transaction completion sent to requestor (SQ de-allocation). See corresponding Umask under OFFCORE_REQUESTS.  |                        |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DATA_RD        | Counts cycles when offcore outstanding cacheable Core Data Read transactions are present in the super queue. A transaction is considered to be in the Offcore outstanding state between L2 miss and transaction completion sent to requestor (SQ de-allocation). See corresponding Umask under OFFCORE_REQUESTS.  | CounterMask=1<br>CMSK1 |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description  | Comment                |
|------------|-------------|---|--|------------------------|
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.L3_MISS_DEMAND_DATA_RD             | Counts number of Offcore outstanding Demand Data Read requests that miss L3 cache in the superQ every cycle.   |                        |
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_L3_MISS_DEMAND_DATA_RD | Cycles with at least 1 Demand Data Read requests who miss L3 cache in the superQ.  | CounterMask=1<br>CMSK1 |
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.L3_MISS_DEMAND_DATA_RD_GE_6        | Cycles with at least 6 Demand Data Read requests that miss L3 cache in the superQ.   | CounterMask=6<br>CMSK6 |
| 79H        | 04H         | IDQ.MITE_UOPS   | Counts the number of uops delivered to Instruction Decode Queue (IDQ) from the MITE path. Counting includes uops that may 'bypass' the IDQ. This also means that uops are not being delivered from the Decode Stream Buffer (DSB).   |                        |
| 79H        | 04H         | IDQ.MITE_CYCLES   | Counts cycles during which uops are being delivered to Instruction Decode Queue (IDQ) from the MITE path. Counting includes uops that may 'bypass' the IDQ.  | CounterMask=1<br>CMSK1 |
| 79H        | 08H         | IDQ.DSB_UOPS  | Counts the number of uops delivered to Instruction Decode Queue (IDQ) from the Decode Stream Buffer (DSB) path. Counting includes uops that may 'bypass' the IDQ.  |                        |
| 79H        | 08H         | IDQ.DSB_CYCLES  | Counts cycles during which uops are being delivered to Instruction Decode Queue (IDQ) from the Decode Stream Buffer (DSB) path. Counting includes uops that may 'bypass' the IDQ.  | CounterMask=1<br>CMSK1 |
| 79H        | 10H         | IDQ.MS_DSB_CYCLES   | Counts cycles during which uops initiated by Decode Stream Buffer (DSB) are being delivered to Instruction Decode Queue (IDQ) while the Microcode Sequencer (MS) is busy. Counting includes uops that may 'bypass' the IDQ.  | CounterMask=1          |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_4_UOPS                                       | Counts the number of cycles 4 uops were delivered to Instruction Decode Queue (IDQ) from the Decode Stream Buffer (DSB) path. Count includes uops that may 'bypass' the IDQ.   | CounterMask=4<br>CMSK4 |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_ANY_UOPS                                     | Counts the number of cycles uops were delivered to Instruction Decode Queue (IDQ) from the Decode Stream Buffer (DSB) path. Count includes uops that may 'bypass' the IDQ.   | CounterMask=1<br>CMSK1 |
| 79H        | 20H         | IDQ.MS_MITE_UOPS  | Counts the number of uops initiated by MITE and delivered to Instruction Decode Queue (IDQ) while the Microcode Sequencer (MS) is busy. Counting includes uops that may 'bypass' the IDQ.  |                        |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_4_UOPS                                      | Counts the number of cycles 4 uops were delivered to the Instruction Decode Queue (IDQ) from the MITE (legacy decode pipeline) path. Counting includes uops that may 'bypass' the IDQ. During these cycles uops are not being delivered from the Decode Stream Buffer (DSB). | CounterMask=4<br>CMSK4 |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description  | Comment                               |
|------------|-------------|-----------------------------------|--|---------------------------------------|
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_ANY_UOPS      | Counts the number of cycles uops were delivered to the Instruction Decode Queue (IDQ) from the MITE (legacy decode pipeline) path. Counting includes uops that may 'bypass' the IDQ. During these cycles uops are not being delivered from the Decode Stream Buffer (DSB). | CounterMask=1<br>CMSK1                |
| 79H        | 30H         | IDQ.MS_CYCLES                     | Counts cycles during which uops are being delivered to Instruction Decode Queue (IDQ) while the Microcode Sequencer (MS) is busy. Counting includes uops that may 'bypass' the IDQ. Uops maybe initiated by Decode Stream Buffer (DSB) or MITE.                            | CounterMask=1<br>CMSK1                |
| 79H        | 30H         | IDQ.MS_SWITCHES                   | Number of switches from DSB (Decode Stream Buffer) or MITE (legacy decode pipeline) to the Microcode Sequencer.  | EdgeDetect=1<br>CounterMask=1<br>EDGE |
| 79H        | 30H         | IDQ.MS_UOPS                       | Counts the total number of uops delivered by the Microcode Sequencer (MS). Any instruction over 4 uops will be delivered by the MS. Some instructions such as transcendentals may additionally generate uops from the MS.  |                                       |
| 80H        | 04H         | ICACHE_16B.IFDATA_STALL           | Cycles where a code line fetch is stalled due to an L1 instruction cache miss. The legacy decode pipeline works at a 16 Byte granularity.  |                                       |
| 83H        | 01H         | ICACHE_64B.IFTAG_HIT              | Instruction fetch tag lookups that hit in the instruction cache (L1). Counts at 64-byte cache-line granularity.  |                                       |
| 83H        | 02H         | ICACHE_64B.IFTAG_MISS             | Instruction fetch tag lookups that miss in the instruction cache (L1). Counts at 64-byte cache-line granularity.   |                                       |
| 83H        | 04H         | ICACHE_64B.IFTAG_STALL            | Cycles where a code fetch is stalled due to L1 instruction cache tag miss.   |                                       |
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK    | Counts page walks of any page size (4K/2M/4M/1G) caused by a code fetch. This implies it missed in the ITLB and further levels of TLB, but the walk need not have completed.   |                                       |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETE_D_4K    | Counts completed page walks (4K page size) caused by a code fetch. This implies it missed in the ITLB and further levels of TLB. The page walk can end with or without a fault.  |                                       |
| 85H        | 04H         | ITLB_MISSES.WALK_COMPLETE_D_2M_4M | Counts completed page walks of any page size (4K/2M/4M/1G) caused by a code fetch. This implies it missed in the ITLB and further levels of TLB. The page walk can end with or without a fault.  |                                       |
| 85H        | 08H         | ITLB_MISSES.WALK_COMPLETE_D_1G    | Counts store misses in all DTLB levels that cause a completed page walk (1G page size). The page walk can end with or without a fault.   |                                       |
| 85H        | 0EH         | ITLB_MISSES.WALK_COMPLETE_D       | Counts completed page walks (2M and 4M page sizes) caused by a code fetch. This implies it missed in the ITLB and further levels of TLB. The page walk can end with or without a fault.  |                                       |



**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                               | Description   | Comment                             |
|------------|-------------|---|---|-------------------------------------|
| 85H        | 10H         | ITLB_MISSES.WALK_PENDING                          | Counts 1 per cycle for each PMH that is busy with a page walk for an instruction fetch request. EPT page walk duration are excluded in Skylake microarchitecture.   |                                     |
| 85H        | 10H         | ITLB_MISSES.WALK_ACTIVE                           | Cycles when at least one PMH is busy with a page walk for code (instruction fetch) request. EPT page walk duration are excluded in Skylake microarchitecture.   | CounterMask=1                       |
| 85H        | 20H         | ITLB_MISSES.STLB_HIT                              | Instruction fetch requests that miss the ITLB and hit the STLB.   |                                     |
| 87H        | 01H         | ILD_STALL.LCP                                     | Counts cycles that the Instruction Length decoder (ILD) stalls occurred due to dynamically changing prefix length of the decoded instruction (by operand size prefix instruction 0x66, address size prefix instruction 0x67 or REX.W for Intel64). Count is proportional to the number of prefixes in a 16B-line. This may result in a three-cycle penalty for each LCP (Length changing prefix) in a 16-byte chunk.  |                                     |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CORE                       | Counts the number of uops not delivered to Resource Allocation Table (RAT) per thread adding "4 - x" when Resource Allocation Table (RAT) is not stalled and Instruction Decode Queue (IDQ) delivers x uops to Resource Allocation Table (RAT) (where x belongs to {0,1,2,3}). Counting does not cover cases when: a. IDQ-Resource Allocation Table (RAT) pipe serves the other thread. b. Resource Allocation Table (RAT) is stalled for the thread (including uop drops and clear BE conditions). c. Instruction Decode Queue (IDQ) delivers four uops. |                                     |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_0_UOPS_DELIV.CORE   | Counts, on the per-thread basis, cycles when no uops are delivered to Resource Allocation Table (RAT). IDQ_Uops_Not_Delivered.core =4.  | CounterMask=4<br>CMSK4              |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_LE_1_UOP_DELIV.CORE | Counts, on the per-thread basis, cycles when less than 1 uop is delivered to Resource Allocation Table (RAT). IDQ_Uops_Not_Delivered.core >= 3.   | CounterMask=3<br>CMSK3              |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_LE_2_UOP_DELIV.CORE | Cycles with less than 2 uops delivered by the front end.  | CounterMask=2<br>CMSK2              |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_LE_3_UOP_DELIV.CORE | Cycles with less than 3 uops delivered by the front end.  | CounterMask=1<br>CMSK1              |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_FE_WAS_OK           | Counts cycles FE delivered 4 uops or Resource Allocation Table (RAT) was stalling FE.   | CounterMask=1<br>Invert=1 CMSK, INV |
| A1H        | 01H         | UOPS_DISPATCHED_PORT.PORT_0                       | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 0.  |                                     |
| A1H        | 02H         | UOPS_DISPATCHED_PORT.PORT_1                       | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 1.  |                                     |
| A1H        | 04H         | UOPS_DISPATCHED_PORT.PORT_2                       | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 2.  |                                     |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment                  |
|------------|-------------|--------------------------------|--|--------------------------|
| A1H        | 08H         | UOPS_DISPATCHED_PORT.PORT_3    | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 3.   |                          |
| A1H        | 10H         | UOPS_DISPATCHED_PORT.PORT_4    | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 4.   |                          |
| A1H        | 20H         | UOPS_DISPATCHED_PORT.PORT_5    | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 5.   |                          |
| A1H        | 40H         | UOPS_DISPATCHED_PORT.PORT_6    | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 6.   |                          |
| A1H        | 80H         | UOPS_DISPATCHED_PORT.PORT_7    | Counts, on the per-thread basis, cycles during which at least one uop is dispatched from the Reservation Station (RS) to port 7.   |                          |
| A2H        | 01H         | RESOURCE_STALLS.ANY            | Counts resource-related stall cycles. Reasons for stalls can be as follows: a. *any* u-arch structure got full (LB, SB, RS, ROB, BOB, LM, Physical Register Reclaim Table (PRRT), or Physical History Table (PHT) slots). b. *any* u-arch structure got empty (like INT/SIMD FreeLists). c. FPU control word (FPCW), MXCSR, and others. This counts cycles that the pipeline back end blocked uop delivery from the front end. |                          |
| A2H        | 08H         | RESOURCE_STALLS.SB             | Counts allocation stall cycles caused by the store buffer (SB) being full. This counts cycles that the pipeline back end blocked uop delivery from the front end.  |                          |
| A3H        | 01H         | CYCLE_ACTIVITY.CYCLES_L2_MISS  | Cycles while L2 cache miss demand load is outstanding.   | CounterMask=1<br>CMSK1   |
| A3H        | 02H         | CYCLE_ACTIVITY.CYCLES_L3_MISS  | Cycles while L3 cache miss demand load is outstanding.   | CounterMask=2<br>CMSK2   |
| A3H        | 04H         | CYCLE_ACTIVITY.STALLS_TOTAL    | Total execution stalls.  | CounterMask=4<br>CMSK4   |
| A3H        | 05H         | CYCLE_ACTIVITY.STALLS_L2_MISS  | Execution stalls while L2 cache miss demand load is outstanding.   | CounterMask=5<br>CMSK5   |
| A3H        | 06H         | CYCLE_ACTIVITY.STALLS_L3_MISS  | Execution stalls while L3 cache miss demand load is outstanding.   | CounterMask=6<br>CMSK6   |
| A3H        | 08H         | CYCLE_ACTIVITY.CYCLES_L1D_MISS | Cycles while L1 cache miss demand load is outstanding.   | CounterMask=8<br>CMSK8   |
| A3H        | 0CH         | CYCLE_ACTIVITY.STALLS_L1D_MISS | Execution stalls while L1 cache miss demand load is outstanding.   | CounterMask=12<br>CMSK12 |
| A3H        | 10H         | CYCLE_ACTIVITY.CYCLES_MEM_ANY  | Cycles while memory subsystem has an outstanding load.   | CounterMask=16<br>CMSK16 |
| A3H        | 14H         | CYCLE_ACTIVITY.STALLS_MEM_ANY  | Execution stalls while memory subsystem has an outstanding load.   | CounterMask=20<br>CMSK20 |
| A6H        | 01H         | EXE_ACTIVITY.EXE_BOUND_0_PORTS | Counts cycles during which no uops were executed on all ports and Reservation Station (RS) was not empty.  |                          |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description  | Comment                |
|------------|-------------|----------------------------------|--|------------------------|
| A6H        | 02H         | EXE_ACTIVITY.1_PORTS_UTIL        | Counts cycles during which a total of 1 uop was executed on all ports and Reservation Station (RS) was not empty.  |                        |
| A6H        | 04H         | EXE_ACTIVITY.2_PORTS_UTIL        | Counts cycles during which a total of 2 uops were executed on all ports and Reservation Station (RS) was not empty.  |                        |
| A6H        | 08H         | EXE_ACTIVITY.3_PORTS_UTIL        | Cycles total of 3 uops are executed on all ports and Reservation Station (RS) was not empty.   |                        |
| A6H        | 10H         | EXE_ACTIVITY.4_PORTS_UTIL        | Cycles total of 4 uops are executed on all ports and Reservation Station (RS) was not empty.   |                        |
| A6H        | 40H         | EXE_ACTIVITY.BOUND_ON_STORES     | Cycles where the Store Buffer was full and no outstanding load.  |                        |
| A8H        | 01H         | LSD.UOPS                         | Number of uops delivered to the back-end by the LSD (Loop Stream Detector).  |                        |
| A8H        | 01H         | LSD.CYCLES_ACTIVE                | Counts the cycles when at least one uop is delivered by the LSD (Loop-stream detector).  | CounterMask=1<br>CMSK1 |
| A8H        | 01H         | LSD.CYCLES_4_UOPS                | Counts the cycles when 4 uops are delivered by the LSD (Loop-stream detector).   | CounterMask=4<br>CMSK4 |
| ABH        | 02H         | DSB2MITE_SWITCHES.PENALTY_CYCLES | Counts Decode Stream Buffer (DSB)-to-MITE switch true penalty cycles. These cycles do not include uops routed through because of the switch itself, for example, when Instruction Decode Queue (IDQ) pre-allocation is unavailable, or Instruction Decode Queue (IDQ) is full. SBD-to-MITE switch true penalty cycles happen after the merge mux (MM) receives Decode Stream Buffer (DSB) Sync-indication until receiving the first MITE uop. MM is placed before Instruction Decode Queue (IDQ) to merge uops being fed from the MITE and Decode Stream Buffer (DSB) paths. Decode Stream Buffer (DSB) inserts the Sync-indication whenever a Decode Stream Buffer (DSB)-to-MITE switch occurs. Penalty: A Decode Stream Buffer (DSB) hit followed by a Decode Stream Buffer (DSB) miss can cost up to six cycles in which no uops are delivered to the IDQ. Most often, such switches from the Decode Stream Buffer (DSB) to the legacy pipeline cost 0 to 2 cycles. |                        |
| AEH        | 01H         | ITLB.ITLB_FLUSH                  | Counts the number of flushes of the big or small ITLB pages. Counting include both TLB Flush (covering all sets) and TLB Set Clear (set-specific).   |                        |
| BOH        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD  | Counts the Demand Data Read requests sent to uncore. Use it in conjunction with OFFCORE_REQUESTS_OUTSTANDING to determine average latency in the uncore.   |                        |
| BOH        | 02H         | OFFCORE_REQUESTS.DEMAND_CODE_RD  | Counts both cacheable and non-cacheable code read requests.  |                        |
| BOH        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO      | Counts the demand RFO (read for ownership) requests including regular RFOs, locks, ItoM.   |                        |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description  | Comment                              |
|------------|-------------|---|--|--------------------------------------|
| B0H        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD            | Counts the demand and prefetch data reads. All Core Data Reads include cacheable 'Demands' and L2 prefetchers (not L3 prefetchers). Counting also covers reads due to page walks resulted from any request type.   |                                      |
| B0H        | 10H         | OFFCORE_REQUESTS.L3_MISS_DEMAND_DATA_RD | Demand Data Read requests who miss L3 cache.   |                                      |
| B0H        | 80H         | OFFCORE_REQUESTS.ALL_REQUESTS           | Counts memory transactions reached the super queue including requests initiated by the core, all L3 prefetches, page walks, etc.   |                                      |
| B1H        | 01H         | UOPS_EXECUTED.THREAD                    | Number of uops to be executed per-thread each cycle.   |                                      |
| B1H        | 01H         | UOPS_EXECUTED.STALL_CYCLES              | Counts cycles during which no uops were dispatched from the Reservation Station (RS) per thread.   | CounterMask=1<br>Invert=1 CMSK, INV  |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_1_UOPS_EXEC     | Cycles where at least 1 uop was executed per-thread.   | CounterMask=1<br>CMSK1               |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_2_UOPS_EXEC     | Cycles where at least 2 uops were executed per-thread.   | CounterMask=2<br>CMSK2               |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_3_UOPS_EXEC     | Cycles where at least 3 uops were executed per-thread.   | CounterMask=3<br>CMSK3               |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_4_UOPS_EXEC     | Cycles where at least 4 uops were executed per-thread.   | CounterMask=4<br>CMSK4               |
| B1H        | 02H         | UOPS_EXECUTED.CORE                      | Number of uops executed from any thread.   |                                      |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_1          | Cycles at least 1 micro-op is executed from any thread on physical core.   | CounterMask=1<br>CMSK1               |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_2          | Cycles at least 2 micro-op is executed from any thread on physical core.   | CounterMask=2<br>CMSK2               |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_3          | Cycles at least 3 micro-op is executed from any thread on physical core.   | CounterMask=3<br>CMSK3               |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_4          | Cycles at least 4 micro-op is executed from any thread on physical core.   | CounterMask=4<br>CMSK4               |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_NONE          | Cycles with no micro-ops executed from any thread on physical core.  | CounterMask=1<br>Invert=1 CMSK1, INV |
| B1H        | 10H         | UOPS_EXECUTED.X87                       | Counts the number of x87 uops executed.  |                                      |
| B2H        | 01H         | OFFCORE_REQUESTS_BUFFER_SQ_FULL         | Counts the number of cases when the offcore requests buffer cannot take more entries for the core. This can happen when the superqueue does not contain eligible entries, or when L1D writeback pending FIFO requests is full. Note: Writeback pending FIFO has six entries. |                                      |
| BDH        | 01H         | TLB_FLUSH.DTLB_THREAD                   | Counts the number of DTLB flush attempts of the thread-specific entries.   |                                      |
| BDH        | 20H         | TLB_FLUSH.STLB_ANY                      | Counts the number of any STLB flush attempts (such as entire, VPID, PCID, InvPage, CR3 write, etc.).   |                                      |
| COH        | 00H         | INST_RETIRED.ANY_P                      | Counts the number of instructions (EOMs) retired. Counting covers macro-fused instructions individually (that is, increments by two).  | See Table 19-1.                      |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment   |
|------------|-------------|--------------------------------|---|---|
| C0H        | 01H         | INST_RETIRED.PREC_DIST         | A version of INST_RETIRED that allows for a more unbiased distribution of samples across instructions retired. It utilizes the Precise Distribution of Instructions Retired (PDIR) feature to mitigate some bias in how retired instructions get sampled. | Precise event capable<br>Requires PEBS on General Counter 1 (PDIR). |
| C1H        | 3FH         | OTHER_ASSISTS.ANY              | Number of times a microcode assist is invoked by HW other than FP-assist. Examples include AD (page Access Dirty) and AVX* related assists.   |   |
| C2H        | 01H         | UOPS_RETIRED.STALL_CYCLES      | This is a non-precise version (that is, does not use PEBS) of the event that counts cycles without actually retired uops.   | CounterMask=1<br>Invert=1 CMSK1, INV                                |
| C2H        | 01H         | UOPS_RETIRED.TOTAL_CYCLES      | Number of cycles using always true condition (uops_ret < 16) applied to non PEBS uops retired event.  | CounterMask=10<br>Invert=1 CMSK10, INV                              |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS      | Counts the retirement slots used.   |   |
| C3H        | 01H         | MACHINE_CLEARS.COUNT           | Number of machine clears (nukes) of any type.   | EdgeDetect=1<br>CounterMask=1<br>CMSK1, EDG                         |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of memory ordering Machine Clears detected. Memory Ordering Machine Clears can result from one of the following: a. memory disambiguation, b. external snoop, or c. cross SMT-HW-thread snoop (stores) hitting load buffer.             |   |
| C3H        | 04H         | MACHINE_CLEARS.SMC             | Counts self-modifying code (SMC) detected, which causes a machine clear.  |   |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts all (macro) branch instructions retired.   | Precise event capable.<br>See Table 19-1.                           |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL    | This is a non-precise version (that is, does not use PEBS) of the event that counts conditional branch instructions retired.  | Precise event capable.<br>PS  |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL      | This is a non-precise version (that is, does not use PEBS) of the event that counts both direct and indirect near call instructions retired.  | Precise event capable.<br>PS  |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETURN    | This is a non-precise version (that is, does not use PEBS) of the event that counts return instructions retired.  | Precise event capable.<br>PS  |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN      | This is a non-precise version (that is, does not use PEBS) of the event that counts not taken branch instructions retired.  |   |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKEN     | This is a non-precise version (that is, does not use PEBS) of the event that counts taken branch instructions retired.  | Precise event capable.<br>PS  |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANCH     | This is a non-precise version (that is, does not use PEBS) of the event that counts far branch instructions retired.  | Precise event capable.<br>PS  |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description   | Comment                                   |
|------------|-------------|-------------------------------------|---|---|
| C5H        | 00H         | BR_MISP_RETIREDD.ALL_BRANC<br>HES   | Counts all the retired branch instructions that were mispredicted by the processor. A branch misprediction occurs when the processor incorrectly predicts the destination of the branch. When the misprediction is discovered at execution, all the instructions executed in the wrong (speculative) path must be discarded, and the processor must start fetching from the correct path. | Precise event capable.<br>See Table 19-1. |
| C5H        | 01H         | BR_MISP_RETIREDD.CONDITIONA<br>L    | This is a non-precise version (that is, does not use PEBS) of the event that counts mispredicted conditional branch instructions retired.   | Precise event capable.<br>PS              |
| C5H        | 02H         | BR_MISP_RETIREDD.NEAR_CALL          | Counts both taken and not taken retired mispredicted direct and indirect near calls, including both register and memory indirect.   | Precise event capable.                    |
| C5H        | 20H         | BR_MISP_RETIREDD.NEAR_TAKE<br>N     | Number of near branch instructions retired that were mispredicted and taken.  | Precise event capable.<br>PS              |
| C6H        | 01H         | FRONTEND_RETIREDD.DSB_MISS          | Counts retired Instructions that experienced DSB (Decode stream buffer, i.e. the decoded instruction-cache) miss.   | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.L1L_MISS          | Retired Instructions who experienced Instruction L1 Cache true miss.  | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.L2L_MISS          | Retired Instructions who experienced Instruction L2 Cache true miss.  | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.ITLB_MISS         | Counts retired Instructions that experienced iTLB (Instruction TLB) true miss.  | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.STLBMIS<br>S      | Counts retired Instructions that experienced STLB (2nd level TLB) true miss.  | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_2  | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 2 cycles which was not interrupted by a back-end stall.   | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_4  | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 4 cycles which was not interrupted by a back-end stall.   | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_8  | Counts retired instructions that are delivered to the back end after a front-end stall of at least 8 cycles. During this period the front end delivered no uops.  | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_16 | Counts retired instructions that are delivered to the back end after a front-end stall of at least 16 cycles. During this period the front end delivered no uops.   | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_32 | Counts retired instructions that are delivered to the back end after a front-end stall of at least 32 cycles. During this period the front end delivered no uops.   | Precise event capable.                    |
| C6H        | 01H         | FRONTEND_RETIREDD.LATENCY_<br>GE_64 | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 64 cycles which was not interrupted by a back-end stall.  | Precise event capable.                    |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                        | Description  | Comment  |
|------------|-------------|--|--|--|
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_128            | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 128 cycles which was not interrupted by a back-end stall.  | Precise event capable.                         |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_256            | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 256 cycles which was not interrupted by a back-end stall.  | Precise event capable.                         |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_512            | Retired instructions that are fetched after an interval where the front end delivered no uops for a period of 512 cycles which was not interrupted by a back-end stall.  | Precise event capable.                         |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_2_BUBBLES_GE_1 | Counts retired instructions that are delivered to the back end after the front end had at least 1 bubble-slot for a period of 2 cycles. A bubble-slot is an empty issue-pipeline slot while there was no RAT stall.  | Precise event capable.                         |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_2_BUBBLES_GE_2 | Retired instructions that are fetched after an interval where the front end had at least 2 bubble-slots for a period of 2 cycles which was not interrupted by a back-end stall.  | Precise event capable.                         |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_GE_2_BUBBLES_GE_3 | Retired instructions that are fetched after an interval where the front end had at least 3 bubble-slots for a period of 2 cycles which was not interrupted by a back-end stall.  | Precise event capable.                         |
| C7H        | 01H         | FP_ARITH_INST_RETIRED.SCALAR_DOUBLE        | Number of SSE/AVX computational scalar double precision floating-point instructions retired. Each count represents 1 computation. Applies to SSE* and AVX* scalar double precision floating-point instructions: ADD SUB MUL DIV MIN MAX SQRT FM(N)ADD/SUB. FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element.                      | Software may treat each count as one DP FLOP.  |
| C7H        | 02H         | FP_ARITH_INST_RETIRED.SCALAR_SINGLE        | Number of SSE/AVX computational scalar single precision floating-point instructions retired. Each count represents 1 computation. Applies to SSE* and AVX* scalar single precision floating-point instructions: ADD SUB MUL DIV MIN MAX RCP RSQRT SQRT FM(N)ADD/SUB. FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element.            | Software may treat each count as one SP FLOP.  |
| C7H        | 04H         | FP_ARITH_INST_RETIRED.128B_PACKED_DOUBLE   | Number of SSE/AVX computational 128-bit packed double precision floating-point instructions retired. Each count represents 2 computations. Applies to SSE* and AVX* packed double precision floating-point instructions: ADD SUB MUL DIV MIN MAX SQRT DPP FM(N)ADD/SUB. DPP and FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element. | Software may treat each count as two DP FLOPs. |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description  | Comment  |
|------------|-------------|--|--|--|
| C7H        | 08H         | FP_ARITH_INST_RETIRED.128B_PACKED_SINGLE | Number of SSE/AVX computational 128-bit packed single precision floating-point instructions retired. Each count represents 4 computations. Applies to SSE* and AVX* packed single precision floating-point instructions: ADD SUB MUL DIV MIN MAX RCP RSQRT SQRT DPP FM(N)ADD/SUB. DPP and FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element. | Software may treat each count as four SP FLOPs.  |
| C7H        | 10H         | FP_ARITH_INST_RETIRED.256B_PACKED_DOUBLE | Number of SSE/AVX computational 256-bit packed double precision floating-point instructions retired. Each count represents 4 computations. Applies to SSE* and AVX* packed double precision floating-point instructions: ADD SUB MUL DIV MIN MAX SQRT DPP FM(N)ADD/SUB. DPP and FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element.           | Software may treat each count as four DP FLOPs.  |
| C7H        | 20H         | FP_ARITH_INST_RETIRED.256B_PACKED_SINGLE | Number of SSE/AVX computational 256-bit packed single precision floating-point instructions retired. Each count represents 8 computations. Applies to SSE* and AVX* packed single precision floating-point instructions: ADD SUB MUL DIV MIN MAX RCP RSQRT SQRT DPP FM(N)ADD/SUB. DPP and FM(N)ADD/SUB instructions count twice as they perform multiple calculations per element. | Software may treat each count as eight SP FLOPs. |
| C7H        | 40H         | FP_ARITH_INST_RETIRED.512B_PACKED_DOUBLE | Number of Packed Double-Precision FP arithmetic instructions (use operation multiplier of 8).  | Only applicable when AVX-512 is enabled.         |
| C7H        | 80H         | FP_ARITH_INST_RETIRED.512B_PACKED_SINGLE | Number of Packed Single-Precision FP arithmetic instructions (use operation multiplier of 16).   | Only applicable when AVX-512 is enabled.         |
| C8H        | 01H         | HLE_RETIRED.START                        | Number of times we entered an HLE region. Does not count nested transactions.  |  |
| C8H        | 02H         | HLE_RETIRED.COMMIT                       | Number of times HLE commit succeeded.  |  |
| C8H        | 04H         | HLE_RETIRED.ABORTED                      | Number of times HLE abort was triggered.   | Precise event capable.                           |
| C8H        | 08H         | HLE_RETIRED.ABORTED_MEM                  | Number of times an HLE execution aborted due to various memory events (e.g., read/write capacity and conflicts).   |  |
| C8H        | 10H         | HLE_RETIRED.ABORTED_TIMER                | Number of times an HLE execution aborted due to hardware timer expiration.   |  |
| C8H        | 20H         | HLE_RETIRED.ABORTED_UNFRIENDLY           | Number of times an HLE execution aborted due to HLE-unfriendly instructions and certain unfriendly events (such as AD assists etc.).   |  |
| C8H        | 40H         | HLE_RETIRED.ABORTED_MEMTYPE              | Number of times an HLE execution aborted due to incompatible memory type.  |  |
| C8H        | 80H         | HLE_RETIRED.ABORTED_EVENTS               | Number of times an HLE execution aborted due to unfriendly events (such as interrupts).  |  |
| C9H        | 01H         | RTM_RETIRED.START                        | Number of times we entered an RTM region. Does not count nested transactions.  |  |
| C9H        | 02H         | RTM_RETIRED.COMMIT                       | Number of times RTM commit succeeded.  |  |
| C9H        | 04H         | RTM_RETIRED.ABORTED                      | Number of times RTM abort was triggered.   | Precise event capable.                           |



**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description  | Comment   |
|------------|-------------|--|--|---|
| C9H        | 08H         | RTM_RETIREDA.BORTED_MEM                | Number of times an RTM execution aborted due to various memory events (e.g. read/write capacity and conflicts).  |   |
| C9H        | 10H         | RTM_RETIREDA.BORTED_TIMER              | Number of times an RTM execution aborted due to uncommon conditions.   |   |
| C9H        | 20H         | RTM_RETIREDA.BORTED_UNFRIENDLY         | Number of times an RTM execution aborted due to HLE-unfriendly instructions.   |   |
| C9H        | 40H         | RTM_RETIREDA.BORTED_MEMTYPE            | Number of times an RTM execution aborted due to incompatible memory type.  |   |
| C9H        | 80H         | RTM_RETIREDA.BORTED_EVENTS             | Number of times an RTM execution aborted due to none of the previous 4 categories (e.g. interrupt).  |   |
| CAH        | 1EH         | FP_ASSIST.ANY                          | Counts cycles with any input and output SSE or x87 FP assist. If an input and output assist are detected on the same cycle the event increments by 1.  | CounterMask=1<br>CMASK1                               |
| CBH        | 01H         | HW_INTERRUPTS.RECEIVED                 | Counts the number of hardware interruptions received by the processor.   |   |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSERTS            | Increments when an entry is added to the Last Branch Record (LBR) array (or removed from the array in case of RETURNS in call stack mode). The event requires LBR enable via IA32_DEBUGCTL MSR and branch type selection via MSR_LBR_SELECT. |   |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_4   | Counts loads when the latency from first dispatch to completion is greater than 4 cycles. Reported latency may be longer than just the memory latency.   | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_8   | Counts loads when the latency from first dispatch to completion is greater than 8 cycles. Reported latency may be longer than just the memory latency.   | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_16  | Counts loads when the latency from first dispatch to completion is greater than 16 cycles. Reported latency may be longer than just the memory latency.  | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_32  | Counts loads when the latency from first dispatch to completion is greater than 32 cycles. Reported latency may be longer than just the memory latency.  | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_64  | Counts loads when the latency from first dispatch to completion is greater than 64 cycles. Reported latency may be longer than just the memory latency.  | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_128 | Counts loads when the latency from first dispatch to completion is greater than 128 cycles. Reported latency may be longer than just the memory latency.   | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_256 | Counts loads when the latency from first dispatch to completion is greater than 256 cycles. Reported latency may be longer than just the memory latency.   | Precise event capable. Specify threshold in MSR 3F6H. |
| CDH        | 01H         | MEM_TRANS_RETIREDA.LOAD_LATENCY_GT_512 | Counts loads when the latency from first dispatch to completion is greater than 512 cycles. Reported latency may be longer than just the memory latency.   | Precise event capable. Specify threshold in MSR 3F6H. |
| DOH        | 11H         | MEM_INST_RETIREDA.STLB_MISS_LOADS      | Retired load instructions that miss the STLB.  | Precise event capable. PSDLA                          |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description  | Comment                      |
|------------|-------------|--------------------------------------|--|------------------------------|
| D0H        | 12H         | MEM_INST_RETIRED.STLB_MISS_STORES    | Retired store instructions that miss the STLB.   | Precise event capable. PSDLA |
| D0H        | 21H         | MEM_INST_RETIRED.LOCK_LOADS          | Retired load instructions with locked access.  | Precise event capable. PSDLA |
| D0H        | 41H         | MEM_INST_RETIRED.SPLIT_LOADS         | Counts retired load instructions that split across a cacheline boundary.   | Precise event capable. PSDLA |
| D0H        | 42H         | MEM_INST_RETIRED.SPLIT_STORES        | Counts retired store instructions that split across a cacheline boundary.  | Precise event capable. PSDLA |
| D0H        | 81H         | MEM_INST_RETIRED.ALL_LOADS           | All retired load instructions.   | Precise event capable. PSDLA |
| D0H        | 82H         | MEM_INST_RETIRED.ALL_STORES          | All retired store instructions.  | Precise event capable. PSDLA |
| D1H        | 01H         | MEM_LOAD_RETIRED.L1_HIT              | Counts retired load instructions with at least one uop that hit in the L1 data cache. This event includes all Sw prefetches and lock instructions regardless of the data source. | Precise event capable. PSDLA |
| D1H        | 02H         | MEM_LOAD_RETIRED.L2_HIT              | Retired load instructions with L2 cache hits as data sources.  | Precise event capable. PSDLA |
| D1H        | 04H         | MEM_LOAD_RETIRED.L3_HIT              | Counts retired load instructions with at least one uop that hit in the L3 cache.   | Precise event capable. PSDLA |
| D1H        | 08H         | MEM_LOAD_RETIRED.L1_MISS             | Counts retired load instructions with at least one uop that missed in the L1 cache.  | Precise event capable. PSDLA |
| D1H        | 10H         | MEM_LOAD_RETIRED.L2_MISS             | Retired load instructions missed L2 cache as data sources.   | Precise event capable. PSDLA |
| D1H        | 20H         | MEM_LOAD_RETIRED.L3_MISS             | Counts retired load instructions with at least one uop that missed in the L3 cache.  | Precise event capable. PSDLA |
| D1H        | 40H         | MEM_LOAD_RETIRED.FB_HIT              | Counts retired load instructions with at least one uop was load missed in L1 but hit FB (Fill Buffers) due to preceding miss to the same cache line with data not ready.         | Precise event capable. PSDLA |
| D2H        | 01H         | MEM_LOAD_L3_HIT_RETIRED.XSNP_MISS    | Retired load instructions which data sources were L3 hit and cross-core snoop missed in on-pkg core cache.   | Precise event capable. PSDLA |
| D2H        | 02H         | MEM_LOAD_L3_HIT_RETIRED.XSNP_HIT     | Retired load instructions which data sources were L3 and cross-core snoop hits in on-pkg core cache.   | Precise event capable. PSDLA |
| D2H        | 04H         | MEM_LOAD_L3_HIT_RETIRED.XSNP_HITM    | Retired load instructions which data sources were HitM responses from shared L3.   | Precise event capable. PSDLA |
| D2H        | 08H         | MEM_LOAD_L3_HIT_RETIRED.XSNP_NONE    | Retired load instructions which data sources were hits in L3 without snoops required.  | Precise event capable. PSDLA |
| D3H        | 01H         | MEM_LOAD_L3_MISS_RETIRED.LOCAL_DRAM  | Retired load instructions which data sources missed L3 but serviced from local DRAM.   | Precise event capable.       |
| D3H        | 02H         | MEM_LOAD_L3_MISS_RETIRED.REMOTE_DRAM | Retired load instructions which data sources missed L3 but serviced from remote dram.  | Precise event capable.       |
| D3H        | 04H         | MEM_LOAD_L3_MISS_RETIRED.REMOTE_HITM | Retired load instructions whose data sources was remote HITM.  | Precise event capable.       |
| D3H        | 08H         | MEM_LOAD_L3_MISS_RETIRED.REMOTE_FWD  | Retired load instructions whose data sources was forwarded from a remote cache.  |                              |

**Table 19-3. Performance Events of the Processor Core Supported in Intel® Xeon® Processor Scalable Family with Skylake Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description  | Comment                |
|------------|-------------|---------------------------|--|------------------------|
| D4H        | 04H         | MEM_LOAD_MISC_RETIRED.UC  | Retired instructions with at least 1 uncacheable load or lock.   | Precise event capable. |
| E6H        | 01H         | BACLEAR.S.ANY             | Counts the number of times the front-end is re-steered when it finds a branch instruction in a fetch line. This occurs for the first time a branch instruction is fetched or when the branch is not tracked by the BPU (Branch Prediction Unit) anymore.   |                        |
| FOH        | 40H         | L2_TRANS.L2_WB            | Counts L2 writebacks that access L2 cache.   |                        |
| F1H        | 1FH         | L2_LINES_IN.ALL           | Counts the number of L2 cache lines filling the L2. Counting does not cover rejects.   |                        |
| F2H        | 01H         | L2_LINES_OUT.SILENT       | Counts the number of lines that are silently dropped by L2 cache when triggered by an L2 cache fill. These lines are typically in Shared state. A non-threaded event.  |                        |
| F2H        | 02H         | L2_LINES_OUT.NON_SILENT   | Counts the number of lines that are evicted by L2 cache when triggered by an L2 cache fill. Those lines can be either in modified state or clean state. Modified lines may either be written back to L3 or directly written to memory and not allocated in L3. Clean lines may either be allocated in L3 or dropped. |                        |
| F2H        | 04H         | L2_LINES_OUT.USELESS_PREF | Counts the number of lines that have been hardware prefetched but not used and now evicted by L2 cache.  |                        |
| F2H        | 04H         | L2_LINES_OUT.USELESS_HWPF | Counts the number of lines that have been hardware prefetched but not used and now evicted by L2 cache.  |                        |
| F4H        | 10H         | SQ_MISC.SPLIT_LOCK        | Counts the number of cache line split locks sent to the uncore.  |                        |
| FEH        | 02H         | IDI_MISC.WB_UPGRADE       | Counts number of cache lines that are allocated and written back to L3 with the intention that they are more likely to be reused shortly.  |                        |
| FEH        | 04H         | IDI_MISC.WB_DOWNGRADE     | Counts number of cache lines that are dropped and not written back to L3 as they are deemed to be less likely to be reused shortly.  |                        |

## 19.3 PERFORMANCE MONITORING EVENTS FOR 6TH GENERATION, 7TH GENERATION AND 8TH GENERATION INTEL® CORE™ PROCESSORS

6th Generation Intel® Core™ processors are based on the Skylake microarchitecture. They support the architectural performance monitoring events listed in Table 19-1. Fixed counters in the core PMU support the architecture events defined in Table 19-2. Model-specific performance monitoring events in the processor core are listed in Table 19-4. The events in Table 19-4 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_4EH and 06\_5EH. Table 19-10 lists performance events supporting Intel TSX (see Section 18.3.6.5) and the events are applicable to processors based on Skylake microarchitecture. Where Skylake microarchitecture implements TSX-related event semantics that differ from Table 19-10, they are listed in Table 19-5.

7th Generation Intel® Core™ processors are based on the Kaby Lake microarchitecture. 8th Generation Intel® Core™ processors are based on the Coffee Lake microarchitecture. Model-specific performance monitoring events in the processor core for these processors are listed in Table 19-4. The events in Table 19-4 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_8EH and 06\_9EH.

The comment column in Table 19-4 uses abbreviated letters to indicate additional conditions applicable to the Event Mask Mnemonic. For event umasks listed in Table 19-4 that do not show “AnyT”, users should refrain from programming “AnyThread =1” in IA32\_PERF\_EVTSELx.

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description   | Comment    |
|------------|-------------|-------------------------------------|---|------------|
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD             | Loads blocked by overlapping with store buffer that cannot be forwarded.  |            |
| 03H        | 08H         | LD_BLOCKS.NO_SR                     | The number of times that split load operations are temporarily blocked because all resources for handling the split accesses are in use.                                  |            |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS     | False dependencies in MOB due to partial compare on address.  |            |
| 08H        | 01H         | DTLB_LOAD_MISSES.MISS_CAUSES_A_WALK | Load misses in all TLB levels that cause a page walk of any page size.  |            |
| 08H        | 0EH         | DTLB_LOAD_MISSES.WALK_COMPLETED     | Load misses in all TLB levels causes a page walk that completes. (All page sizes.)  |            |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_PENDING       | Counts 1 per cycle for each PMH that is busy with a page walk for a load.   |            |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_ACTIVE        | Cycles when at least one PMH is busy with a walk for a load.  | CMSK1      |
| 08H        | 20H         | DTLB_LOAD_MISSES.STLB_HIT           | Loads that miss the DTLB but hit STLB.  |            |
| 0DH        | 01H         | INT_MISC.RECOVERY_CYCLES            | Core cycles the allocator was stalled due to recovery from earlier machine clear event for this thread (for example, misprediction or memory order conflict).             |            |
| 0DH        | 01H         | INT_MISC.RECOVERY_CYCLES_ANY        | Core cycles the allocator was stalled due to recovery from earlier machine clear event for any logical thread in this processor core.                                     | AnyT       |
| 0DH        | 80H         | INT_MISC.CLEAR_RESTEER_CYCLES       | Cycles the issue-stage is waiting for front end to fetch from resteeered path following branch misprediction or machine clear events.                                     |            |
| 0EH        | 01H         | UOPS_ISSUED.ANY                     | The number of uops issued by the RAT to RS.   |            |
| 0EH        | 01H         | UOPS_ISSUED.STALL_CYCLES            | Cycles when the RAT does not issue uops to RS for the thread.   | CMSK1, INV |
| 0EH        | 02H         | UOPS_ISSUED.VECTOR_WIDTH_MISMATCH   | Uops inserted at issue-stage in order to preserve upper bits of vector registers.   |            |
| 0EH        | 20H         | UOPS_ISSUED.SLOW_LEA                | Number of slow LEA or similar uops allocated. Such uop has 3 sources (for example, 2 sources + immediate) regardless of whether it is a result of LEA instruction or not. |            |
| 14H        | 01H         | ARITH.FPU_DIVIDER_ACTIVE            | Cycles when divider is busy executing divide or square root operations. Accounts for FP operations including integer divides.   |            |
| 24H        | 21H         | L2_RQSTS.DEMAND_DATA_RD_MISS        | Demand Data Read requests that missed L2, no rejects.   |            |
| 24H        | 22H         | L2_RQSTS.RFO_MISS                   | RFO requests that missed L2.  |            |
| 24H        | 24H         | L2_RQSTS.CODE_RD_MISS               | L2 cache misses when fetching instructions.   |            |
| 24H        | 27H         | L2_RQSTS.ALL_DEMAND_MISS            | Demand requests that missed L2.   |            |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description   | Comment         |
|------------|-------------|---|---|-----------------|
| 24H        | 38H         | L2_RQSTS.PF_MISS                          | Requests from the L1/L2/L3 hardware prefetchers or load software prefetches that miss L2 cache.   |                 |
| 24H        | 3FH         | L2_RQSTS.MISS                             | All requests that missed L2.  |                 |
| 24H        | 41H         | L2_RQSTS.DEMAND_DATA_RD_HIT               | Demand Data Read requests that hit L2 cache.  |                 |
| 24H        | 42H         | L2_RQSTS.RFO_HIT                          | RFO requests that hit L2 cache.   |                 |
| 24H        | 44H         | L2_RQSTS.CODE_RD_HIT                      | L2 cache hits when fetching instructions.   |                 |
| 24H        | D8H         | L2_RQSTS.PF_HIT                           | Prefetches that hit L2.   |                 |
| 24H        | E1H         | L2_RQSTS.ALL_DEMAND_DATA_RD               | All demand data read requests to L2.  |                 |
| 24H        | E2H         | L2_RQSTS.ALL_RFO                          | All L RFO requests to L2.   |                 |
| 24H        | E4H         | L2_RQSTS.ALL_CODE_RD                      | All L2 code requests.   |                 |
| 24H        | E7H         | L2_RQSTS.ALL_DEMAND_REFERENCES            | All demand requests to L2.  |                 |
| 24H        | F8H         | L2_RQSTS.ALL_PF                           | All requests from the L1/L2/L3 hardware prefetchers or load software prefetches.  |                 |
| 24H        | EFH         | L2_RQSTS.REFERENCES                       | All requests to L2.   |                 |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE               | This event counts requests originating from the core that reference a cache line in the L3 cache.   | See Table 19-1. |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS                    | This event counts each cache miss condition for references to the L3 cache.   | See Table 19-1. |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P                 | Cycles while the logical processor is not in a halt state.  | See Table 19-1. |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P_ANY             | Cycles while at least one logical processor is not in a halt state.   | AnyT            |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK          | Core crystal clock cycles when the thread is unhalted.  | See Table 19-1. |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK_ANY      | Core crystal clock cycles when at least one thread on the physical core is unhalted.  | AnyT            |
| 3CH        | 02H         | CPU_CLK_THREAD_UNHALTED.ONE_THREAD_ACTIVE | Core crystal clock cycles when this thread is unhalted and the other thread is halted.  |                 |
| 48H        | 01H         | L1D_PEND_MISS.PENDING                     | Increments the number of outstanding L1D misses every cycle.  |                 |
| 48H        | 01H         | L1D_PEND_MISS.PENDING_CYCLES              | Cycles with at least one outstanding L1D misses from this logical processor.  | CMSK1           |
| 48H        | 01H         | L1D_PEND_MISS.PENDING_CYCLES_ANY          | Cycles with at least one outstanding L1D misses from any logical processor in this core.  | CMSK1, AnyT     |
| 48H        | 02H         | L1D_PEND_MISS.FB_FULL                     | Number of times a request needed a FB entry but there was no entry available for it. That is, the FB unavailability was the dominant reason for blocking the request. A request includes cacheable/uncacheable demand that is load, store or SW prefetch. HWP are excluded. |                 |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK      | Store misses in all TLB levels that cause page walks.   |                 |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                                     | Description  | Comment    |
|------------|-------------|---|--|------------|
| 49H        | 0EH         | DTLB_STORE_MISSES.WALK_COMPLETED                        | Counts completed page walks in any TLB levels due to store misses (all page sizes).  |            |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_PENDING                          | Counts 1 per cycle for each PMH that is busy with a page walk for a store.   |            |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_ACTIVE                           | Cycles when at least one PMH is busy with a page walk for a store.   | CMSK1      |
| 49H        | 20H         | DTLB_STORE_MISSES.STLB_HIT                              | Store misses that missed DTLB but hit STLB.  |            |
| 4CH        | 01H         | LOAD_HIT_PRE.HW_PF                                      | Demand load dispatches that hit fill buffer allocated for software prefetch.   |            |
| 4FH        | 10H         | EPT.WALK_PENDING  | Counts 1 per cycle for each PMH that is busy with an EPT walk for any request type.  |            |
| 51H        | 01H         | L1D.REPLACEMENT   | Counts the number of lines brought into the L1 data cache.   |            |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES                                  | Cycles the RS is empty for the thread.   |            |
| 5EH        | 01H         | RS_EVENTS.EMPTY_END                                     | Counts end of periods where the Reservation Station (RS) was empty. Could be useful to precisely locate Front-end Latency Bound issues.  | CMSK1, INV |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD             | Increment each cycle of the number of offcore outstanding Demand Data Read transactions in SQ to uncure.                                 |            |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_DATA_RD | Cycles with at least one offcore outstanding Demand Data Read transactions in SQ to uncure.  | CMSK1      |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD_GE_6        | Cycles with at least 6 offcore outstanding Demand Data Read transactions in SQ to uncure.  | CMSK6      |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_CODE_RD             | Increment each cycle of the number of offcore outstanding demand code read transactions in SQ to uncure.                                 |            |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_CODE_RD | Cycles with at least one offcore outstanding demand code read transactions in SQ to uncure.  | CMSK1      |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO                 | Increment each cycle of the number of offcore outstanding RFO store transactions in SQ to uncure. Set Cmask=1 to count cycles.           |            |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DEMAND_RFO     | Cycles with at least one offcore outstanding RFO transactions in SQ to uncure.   | CMSK1      |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD                | Increment each cycle of the number of offcore outstanding cacheable data read transactions in SQ to uncure. Set Cmask=1 to count cycles. |            |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_DATA_RD        | Cycles with at least one offcore outstanding data read transactions in SQ to uncure.   | CMSK1      |
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.L3_MISS_DEMAND_DATA_RD     | Increment each cycle of the number of offcore outstanding demand data read requests from SQ that missed L3.                              |            |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description  | Comment |
|------------|-------------|---|--|---------|
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.CYCLES_WITH_L3_MISS_DEMAND_DATA_RD | Cycles with at least one offcore outstanding demand data read requests from SQ that missed L3.                   | CMSK1   |
| 60H        | 10H         | OFFCORE_REQUESTS_OUTSTANDING.L3_MISS_DEMAND_DATA_RD_GE_6        | Cycles with at least one offcore outstanding demand data read requests from SQ that missed L3.                   | CMSK6   |
| 63H        | 02H         | LOCK_CYCLES.CACHE_LOCK_DURATION                                 | Cycles in which the L1D is locked.   |         |
| 79H        | 04H         | IDQ.MITE_UOPS   | Increment each cycle # of uops delivered to IDQ from MITE path.  |         |
| 79H        | 04H         | IDQ.MITE_CYCLES   | Cycles when uops are being delivered to IDQ from MITE path.  | CMSK1   |
| 79H        | 08H         | IDQ.DSB_UOPS  | Increment each cycle. # of uops delivered to IDQ from DSB path.  |         |
| 79H        | 08H         | IDQ.DSB_CYCLES  | Cycles when uops are being delivered to IDQ from DSB path.   | CMSK1   |
| 79H        | 10H         | IDQ.MS_DSB_UOPS   | Increment each cycle # of uops delivered to IDQ by DSB when MS_busy.   |         |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_ANY_UOPS                                     | Cycles DSB is delivered at least one uops.   | CMSK1   |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_4_UOPS                                       | Cycles DSB is delivered four uops.   | CMSK4   |
| 79H        | 20H         | IDQ.MS_MITE_UOPS  | Increment each cycle # of uops delivered to IDQ by MITE when MS_busy.  |         |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_ANY_UOPS                                    | Counts cycles MITE is delivered at least one uops.   | CMSK1   |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_4_UOPS                                      | Counts cycles MITE is delivered four uops.   | CMSK4   |
| 79H        | 30H         | IDQ.MS_UOPS   | Increment each cycle # of uops delivered to IDQ while MS is busy.  |         |
| 79H        | 30H         | IDQ.MS_SWITCHES   | Number of switches from DSB or MITE to MS.   | EDG     |
| 79H        | 30H         | IDQ.MS_CYCLES   | Cycles MS is delivered at least one uops.  | CMSK1   |
| 80H        | 04H         | ICACHE_16B.IFDATA_STALL   | Cycles where a code fetch is stalled due to L1 instruction cache miss.   |         |
| 80H        | 04H         | ICACHE_64B.IFDATA_STALL   | Cycles where a code fetch is stalled due to L1 instruction cache tag miss.                                       |         |
| 83H        | 01H         | ICACHE_64B.IFTAG_HIT  | Instruction fetch tag lookups that hit in the instruction cache (L1). Counts at 64-byte cache-line granularity.  |         |
| 83H        | 02H         | ICACHE_64B.IFTAG_MISS   | Instruction fetch tag lookups that miss in the instruction cache (L1). Counts at 64-byte cache-line granularity. |         |
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK                                  | Misses at all ITLB levels that cause page walks.   |         |
| 85H        | 0EH         | ITLB_MISSES.WALK_COMPLETED                                      | Counts completed page walks in any TLB level due to code fetch misses (all page sizes).                          |         |
| 85H        | 10H         | ITLB_MISSES.WALK_PENDING  | Counts 1 per cycle for each PMH that is busy with a page walk for an instruction fetch request.                  |         |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                               | Description  | Comment                     |
|------------|-------------|---|--|-----------------------------|
| 85H        | 20H         | ITLB_MISSES.STLB_HIT                              | ITLB misses that hit STLB.   |                             |
| 87H        | 01H         | ILD_STALL.LCP                                     | Stalls caused by changing prefix length of the instruction.  |                             |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CORE                       | Count issue pipeline slots where no uop was delivered from the front end to the back end when there is no back-end stall.        |                             |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_0_UOP_DELIV.CORE    | Cycles which 4 issue pipeline slots had no uop delivered from the front end to the back end when there is no back-end stall.     | CMSK4                       |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_LE_n_UOP_DELIV.CORE | Cycles which "4-n" issue pipeline slots had no uop delivered from the front end to the back end when there is no back-end stall. | Set CMSK = 4-n; n = 1, 2, 3 |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CYCLES_FE_WAS_OK           | Cycles which front end delivered 4 uops or the RAT was stalling FE.  | CMSK, INV                   |
| A1H        | 01H         | UOPS_DISPATCHED_PORT.PORT_0                       | Counts the number of cycles in which a uop is dispatched to port 0.  |                             |
| A1H        | 02H         | UOPS_DISPATCHED_PORT.PORT_1                       | Counts the number of cycles in which a uop is dispatched to port 1.  |                             |
| A1H        | 04H         | UOPS_DISPATCHED_PORT.PORT_2                       | Counts the number of cycles in which a uop is dispatched to port 2.  |                             |
| A1H        | 08H         | UOPS_DISPATCHED_PORT.PORT_3                       | Counts the number of cycles in which a uop is dispatched to port 3.  |                             |
| A1H        | 10H         | UOPS_DISPATCHED_PORT.PORT_4                       | Counts the number of cycles in which a uop is dispatched to port 4.  |                             |
| A1H        | 20H         | UOPS_DISPATCHED_PORT.PORT_5                       | Counts the number of cycles in which a uop is dispatched to port 5.  |                             |
| A1H        | 40H         | UOPS_DISPATCHED_PORT.PORT_6                       | Counts the number of cycles in which a uop is dispatched to port 6.  |                             |
| A1H        | 80H         | UOPS_DISPATCHED_PORT.PORT_7                       | Counts the number of cycles in which a uop is dispatched to port 7.  |                             |
| A2H        | 01H         | RESOURCE_STALLS.ANY                               | Resource-related stall cycles.   |                             |
| A2H        | 08H         | RESOURCE_STALLS.SB                                | Cycles stalled due to no store buffers available (not including draining from sync).   |                             |
| A3H        | 01H         | CYCLE_ACTIVITY.CYCLES_L2_MISS                     | Cycles while L2 cache miss demand load is outstanding.   | CMSK1                       |
| A3H        | 02H         | CYCLE_ACTIVITY.CYCLES_L3_MISS                     | Cycles while L3 cache miss demand load is outstanding.   | CMSK2                       |
| A3H        | 04H         | CYCLE_ACTIVITY.STALLS_TOTAL                       | Total execution stalls.  | CMSK4                       |
| A3H        | 05H         | CYCLE_ACTIVITY.STALLS_L2_MISS                     | Execution stalls while L2 cache miss demand load is outstanding.   | CMSK5                       |
| A3H        | 06H         | CYCLE_ACTIVITY.STALLS_L3_MISS                     | Execution stalls while L3 cache miss demand load is outstanding.   | CMSK6                       |
| A3H        | 08H         | CYCLE_ACTIVITY.CYCLES_L1D_MISS                    | Cycles while L1 data cache miss demand load is outstanding.  | CMSK8                       |
| A3H        | 0CH         | CYCLE_ACTIVITY.STALLS_L1D_MISS                    | Execution stalls while L1 data cache miss demand load is outstanding.  | CMSK12                      |



**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment   |
|------------|-------------|---|---|-----------|
| A3H        | 10H         | CYCLE_ACTIVITY.CYCLES_MEM_ANY           | Cycles while memory subsystem has an outstanding load.  | CMSK16    |
| A3H        | 14H         | CYCLE_ACTIVITY.STALLS_MEM_ANY           | Execution stalls while memory subsystem has an outstanding load.  | CMSK20    |
| A6H        | 01H         | EXE_ACTIVITY.EXE_BOUND_0_PORTS          | Cycles for which no uops began execution, the Reservation Station was not empty, the Store Buffer was full and there was no outstanding load. |           |
| A6H        | 02H         | EXE_ACTIVITY.1_PORTS_UTIL               | Cycles for which one uop began execution on any port, and the Reservation Station was not empty.  |           |
| A6H        | 04H         | EXE_ACTIVITY.2_PORTS_UTIL               | Cycles for which two uops began execution, and the Reservation Station was not empty.   |           |
| A6H        | 08H         | EXE_ACTIVITY.3_PORTS_UTIL               | Cycles for which three uops began execution, and the Reservation Station was not empty.   |           |
| A6H        | 04H         | EXE_ACTIVITY.4_PORTS_UTIL               | Cycles for which four uops began execution, and the Reservation Station was not empty.  |           |
| A6H        | 40H         | EXE_ACTIVITY.BOUND_ON_STORES            | Cycles where the Store Buffer was full and no outstanding load.   |           |
| A8H        | 01H         | LSD.UOPS                                | Number of uops delivered by the LSD.  |           |
| A8H        | 01H         | LSD.CYCLES_ACTIVE                       | Cycles with at least one uop delivered by the LSD and none from the decoder.  | CMSK1     |
| A8H        | 01H         | LSD.CYCLES_4_UOPS                       | Cycles with 4 uops delivered by the LSD and none from the decoder.  | CMSK4     |
| ABH        | 02H         | DSB2MITE_SWITCHES.PENALTY_CYCLES        | DSB-to-MITE switch true penalty cycles.   |           |
| AEH        | 01H         | ITLB.ITLB_FLUSH                         | Flushing of the Instruction TLB (ITLB) pages, includes 4k/2M/4M pages.  |           |
| B0H        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD         | Demand data read requests sent to uncore.   |           |
| B0H        | 02H         | OFFCORE_REQUESTS.DEMAND_CODE_RD         | Demand code read requests sent to uncore.   |           |
| B0H        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO             | Demand RFO read requests sent to uncore, including regular RFOs, locks, ItoM.   |           |
| B0H        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD            | Data read requests sent to uncore (demand and prefetch).  |           |
| B0H        | 10H         | OFFCORE_REQUESTS.L3_MISS_DEMAND_DATA_RD | Demand data read requests that missed L3.   |           |
| B0H        | 80H         | OFFCORE_REQUESTS.ALL_REQUESTS           | Any memory transaction that reached the SQ.   |           |
| B1H        | 01H         | UOPS_EXECUTED.THREAD                    | Counts the number of uops that begin execution across all ports.  |           |
| B1H        | 01H         | UOPS_EXECUTED.STALL_CYCLES              | Cycles where there were no uops that began execution.   | CMSK, INV |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_1_UOP_EXEC      | Cycles where there was at least one uop that began execution.   | CMSK1     |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_2_UOP_EXEC      | Cycles where there were at least two uops that began execution.   | CMSK2     |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                | Description   | Comment            |
|------------|-------------|------------------------------------|---|--------------------|
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_3_UOP_EXEC | Cycles where there were at least three uops that began execution.   | CMSK3              |
| B1H        | 01H         | UOPS_EXECUTED.CYCLES_GE_4_UOP_EXEC | Cycles where there were at least four uops that began execution.  | CMSK4              |
| B1H        | 02H         | UOPS_EXECUTED.CORE                 | Counts the number of uops from any logical processor in this core that begin execution.   |                    |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_1     | Cycles where there was at least one uop, from any logical processor in this core, that began execution.                                     | CMSK1              |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_2     | Cycles where there were at least two uops, from any logical processor in this core, that began execution.                                   | CMSK2              |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_3     | Cycles where there were at least three uops, from any logical processor in this core, that began execution.                                 | CMSK3              |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_GE_4     | Cycles where there were at least four uops, from any logical processor in this core, that began execution.                                  | CMSK4              |
| B1H        | 02H         | UOPS_EXECUTED.CORE_CYCLES_NONE     | Cycles where there were no uops from any logical processor in this core that began execution.   | CMSK1, INV         |
| B1H        | 10H         | UOPS_EXECUTED.X87                  | Counts the number of X87 uops that begin execution.   |                    |
| B2H        | 01H         | OFF_CORE_REQUEST_BUFFER.SQ_FULL    | Offcore requests buffer cannot take more entries for this core.   |                    |
| B7H        | 01H         | OFF_CORE_RESPONSE_0                | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A6H |
| BBH        | 01H         | OFF_CORE_RESPONSE_1                | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A7H |
| BDH        | 01H         | TLB_FLUSH.DTLB_THREAD              | DTLB flush attempts of the thread-specific entries.   |                    |
| BDH        | 01H         | TLB_FLUSH.STLB_ANY                 | STLB flush attempts.  |                    |
| COH        | 00H         | INST_RETIRED.ANY_P                 | Number of instructions at retirement.   | See Table 19-1.    |
| COH        | 01H         | INST_RETIRED.PREC_DIST             | Precise instruction retired event with HW to reduce effect of PEBS shadow in IP distribution.   | PMC1 only;         |
| COH        | 01H         | INST_RETIRED.TOTAL_CYCLES          | Number of cycles using always true condition applied to PEBS instructions retired event.  | CMSK10, PS         |
| C1H        | 3FH         | OTHER_ASSISTS.ANY                  | Number of times a microcode assist is invoked by HW other than FP-assist. Examples include AD (page Access Dirty) and AVX* related assists. |                    |
| C2H        | 01H         | UOPS_RETIRED.STALL_CYCLES          | Cycles without actually retired uops.   | CMSK1, INV         |
| C2H        | 01H         | UOPS_RETIRED.TOTAL_CYCLES          | Cycles with less than 10 actually retired uops.   | CMSK10, INV        |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS          | Retirement slots used.  |                    |
| C3H        | 01H         | MACHINE_CLEARS.COUNT               | Number of machine clears of any type.   | CMSK1, EDG         |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING     | Counts the number of machine clears due to memory order conflicts.  |                    |
| C3H        | 04H         | MACHINE_CLEARS.SMC                 | Number of self-modifying-code machine clears detected.  |                    |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES       | Branch instructions that retired.   | See Table 19-1.    |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL        | Counts the number of conditional branch instructions retired.   | PS                 |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                            | Description  | Comment                                       |
|------------|-------------|--|--|---|
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL                      | Direct and indirect near call instructions retired.  | PS  |
| C4H        | 04H         | BR_INST_RETIRED.ALL_BRANC<br>HES               | Counts the number of branch instructions retired.  | PS  |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETU<br>RN                | Counts the number of near return instructions retired.   | PS  |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN                      | Counts the number of not taken branch instructions retired.  |   |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKE<br>N                 | Number of near taken branches retired.   | PS  |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANC<br>H                 | Number of far branches retired.  | PS  |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANC<br>HES               | Mispredicted branch instructions at retirement.  | See Table 19-1.                               |
| C5H        | 01H         | BR_MISP_RETIRED.CONDITIONA<br>L                | Mispredicted conditional branch instructions retired.  | PS  |
| C5H        | 04H         | BR_MISP_RETIRED.ALL_BRANC<br>HES               | Mispredicted macro branch instructions retired.  | PS  |
| C5H        | 20H         | BR_MISP_RETIRED.NEAR_TAKE<br>N                 | Number of near branch instructions retired that were mispredicted and taken.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.DSB_MISS                      | Retired instructions which experienced DSB miss. Specify MSR_PEBS_FRONTEND.EVTSEL=11H.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.L1I_MISS                      | Retired instructions which experienced instruction L1 cache true miss. Specify MSR_PEBS_FRONTEND.EVTSEL=12H.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.L2_MISS                       | Retired instructions which experienced L2 cache true miss. Specify MSR_PEBS_FRONTEND.EVTSEL=13H.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.ITLB_MISS                     | Retired instructions which experienced ITLB true miss. Specify MSR_PEBS_FRONTEND.EVTSEL=14H.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.STLB_MIS<br>S                 | Retired instructions which experienced STLB true miss. Specify MSR_PEBS_FRONTEND.EVTSEL=15H.   | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_<br>GE_16             | Retired instructions that are fetched after an interval where the front end delivered no uops for at least 16 cycles. Specify the following fields in MSR_PEBS_FRONTEND: EVTSEL=16H, IDQ_Bubble_Length =16, IDQ_Bubble_Width = 4.                  | PS  |
| C6H        | 01H         | FRONTEND_RETIRED.LATENCY_<br>GE_2_BUBBLES_GE_m | Retired instructions that are fetched after an interval where the front end had 'm' IDQ slots delivered, no uops for at least 2 cycles. Specify the following fields in MSR_PEBS_FRONTEND: EVTSEL=16H, IDQ_Bubble_Length =2, IDQ_Bubble_Width = m. | PS, m = 1, 2, 3                               |
| C7H        | 01H         | FP_ARITH_INST_RETIRED.SCAL<br>AR_DOUBLE        | Number of double-precision, floating-point, scalar SSE/AVX computational instructions that are retired. Each scalar FMA instruction counts as 2.   | Software may treat each count as one DP FLOP. |
| C7H        | 02H         | FP_ARITH_INST_RETIRED.SCAL<br>AR_SINGLE        | Number of single-precision, floating-point, scalar SSE/AVX computational instructions that are retired. Each scalar FMA instruction counts as 2.   | Software may treat each count as one SP FLOP. |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description  | Comment  |
|------------|-------------|--|--|--|
| C7H        | 04H         | FP_ARITH_INST_RETIRED.128B_PACKED_DOUBLE | Number of double-precision, floating-point, 128-bit SSE/AVX computational instructions that are retired. Each 128-bit FMA or (V)DPPD instruction counts as 2.  | Software may treat each count as two DP FLOPs.   |
| C7H        | 08H         | FP_ARITH_INST_RETIRED.128B_PACKED_SINGLE | Number of single-precision, floating-point, 128-bit SSE/AVX computational instructions that are retired. Each 128-bit FMA or (V)DPPS instruction counts as 2.  | Software may treat each count as four SP FLOPs.  |
| C7H        | 10H         | FP_ARITH_INST_RETIRED.256B_PACKED_DOUBLE | Number of double-precision, floating-point, 256-bit SSE/AVX computational instructions that are retired. Each 256-bit FMA instruction counts as 2.   | Software may treat each count as four DP FLOPs.  |
| C7H        | 20H         | FP_ARITH_INST_RETIRED.256B_PACKED_SINGLE | Number of single-precision, floating-point, 256-bit SSE/AVX computational instructions that are retired. Each 256-bit FMA or VDPPS instruction counts as 2.  | Software may treat each count as eight SP FLOPs. |
| CAH        | 1EH         | FP_ASSIST.ANY                            | Cycles with any input/output SSE* or FP assists.   | CMSK1  |
| CBH        | 01H         | HW_INTERRUPTS.RECEIVED                   | Number of hardware interrupts received by the processor.   |  |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSERTS              | Increments when an entry is added to the Last Branch Record (LBR) array (or removed from the array in case of RETURNS in call stack mode). The event requires LBR enable via IA32_DEBUGCTL MSR and branch type selection via MSR_LBR_SELECT. |  |
| CDH        | 01H         | MEM_TRANS_RETIRED.LOAD_LATENCY           | Randomly sampled loads whose latency is above a user defined threshold. A small fraction of the overall loads are sampled due to randomization.  | Specify threshold in MSR 3F6H.<br>PSDLA          |
| DOH        | 11H         | MEM_INST_RETIRED.STLB_MISS_LOADS         | Retired load instructions that miss the STLB.  | PSDLA  |
| DOH        | 12H         | MEM_INST_RETIRED.STLB_MISS_STORES        | Retired store instructions that miss the STLB.   | PSDLA  |
| DOH        | 21H         | MEM_INST_RETIRED.LOCK_LOADS              | Retired load instructions with locked access.  | PSDLA  |
| DOH        | 41H         | MEM_INST_RETIRED.SPLIT_LOADS             | Number of load instructions retired with cache-line splits that may impact performance.  | PSDLA  |
| DOH        | 42H         | MEM_INST_RETIRED.SPLIT_STORES            | Number of store instructions retired with line-split.  | PSDLA  |
| DOH        | 81H         | MEM_INST_RETIRED.ALL_LOADS               | All retired load instructions.   | PSDLA  |
| DOH        | 82H         | MEM_INST_RETIRED.ALL_STORES              | All retired store instructions.  | PSDLA  |
| D1H        | 01H         | MEM_LOAD_RETIRED.L1_HIT                  | Retired load instructions with L1 cache hits as data sources.  | PSDLA  |
| D1H        | 02H         | MEM_LOAD_RETIRED.L2_HIT                  | Retired load instructions with L2 cache hits as data sources.  | PSDLA  |
| D1H        | 04H         | MEM_LOAD_RETIRED.L3_HIT                  | Retired load instructions with L3 cache hits as data sources.  | PSDLA  |
| D1H        | 08H         | MEM_LOAD_RETIRED.L1_MISS                 | Retired load instructions missed L1 cache as data sources.   | PSDLA  |
| D1H        | 10H         | MEM_LOAD_RETIRED.L2_MISS                 | Retired load instructions missed L2. Unknown data source excluded.   | PSDLA  |

**Table 19-4. Performance Events of the Processor Core Supported by Skylake, Kaby Lake and Coffee Lake Microarchitectures (Contd.)**

| Event Num.  | Umask Value | Event Mask Mnemonic                | Description  | Comment |
|---|-------------|------------------------------------|--|---------|
| D1H   | 20H         | MEM_LOAD_RETIRED.L3_MISS           | Retired load instructions missed L3. Excludes unknown data source.   | PSDLA   |
| D1H   | 40H         | MEM_LOAD_RETIRED.FB_HIT            | Retired load instructions where data sources were load uops missed L1 but hit FB due to preceding miss to the same cache line with data not ready. | PSDLA   |
| D2H   | 01H         | MEM_LOAD_L3_HIT_RETIRED.X_SNP_MISS | Retired load instructions where data sources were L3 hit and cross-core snoop missed in on-pkg core cache.   | PSDLA   |
| D2H   | 02H         | MEM_LOAD_L3_HIT_RETIRED.X_SNP_HIT  | Retired load Instructions where data sources were L3 and cross-core snoop hits in on-pkg core cache.   | PSDLA   |
| D2H   | 04H         | MEM_LOAD_L3_HIT_RETIRED.X_SNP_HITM | Retired load instructions where data sources were HitM responses from shared L3.   | PSDLA   |
| D2H   | 08H         | MEM_LOAD_L3_HIT_RETIRED.X_SNP_NONE | Retired load instructions where data sources were hits in L3 without snoops required.  | PSDLA   |
| E6H   | 01H         | BACLEAR.S.ANY                      | Number of front end re-steers due to BPU misprediction.  |         |
| FOH   | 40H         | L2_TRANS.L2_WB                     | L2 writebacks that access L2 cache.  |         |
| F1H   | 07H         | L2_LINES_IN.ALL                    | L2 cache lines filling L2.   |         |
| CMSK1: Counter Mask = 1 required; CMSK4: CounterMask = 4 required; CMSK6: CounterMask = 6 required; CMSK8: CounterMask = 8 required; CMSK10: CounterMask = 10 required; CMSK12: CounterMask = 12 required; CMSK16: CounterMask = 16 required; CMSK20: CounterMask = 20 required.<br>AnyT: AnyThread = 1 required.<br>INV: Invert = 1 required.<br>EDG: EDGE = 1 required.<br>PSDLA: Also supports PEBS and DataLA.<br>PS: Also supports PEBS. |             |                                    |  |         |

Table 19-10 lists performance events supporting Intel TSX (see Section 18.3.6.5) and the events are applicable to processors based on Skylake microarchitecture. Where Skylake microarchitecture implements TSX-related event semantics that differ from Table 19-10, they are listed in Table 19-5.

**Table 19-5. Intel® TSX Performance Event Addendum in Processors based on Skylake Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description   | Comment |
|------------|-------------|-----------------------|---|---------|
| 54H        | 02H         | TX_MEM.ABORT_CAPACITY | Number of times a transactional abort was signaled due to a data capacity limitation for transactional reads or writes. |         |

## 19.4 PERFORMANCE MONITORING EVENTS FOR INTEL® XEON PHI™ PROCESSOR 3200, 5200, 7200 SERIES AND INTEL® XEON PHI™ PROCESSOR 7215, 7285, 7295 SERIES

The Intel® Xeon Phi™ processor 3200/5200/7200 series is based on the Knights Landing microarchitecture with CPUID DisplayFamily\_DisplayModel signature 06\_57H. The Intel® Xeon Phi™ processor 7215/7285/7295 series is based on the Knights Mill microarchitecture with CPUID DisplayFamily\_DisplayModel signature 06\_85H. Model-specific performance monitoring events in these processor cores are listed in Table 19-6. The events in Table 19-6

apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_57H and 06\_85H.

**Table 19-6. Performance Events of the Processor Core Supported by Knights Landing and Knights Mill Microarchitectures**

| Event Num. | Umask Value | Event Mask Mnemonic                | Description  | Comment      |
|------------|-------------|------------------------------------|--|--------------|
| 03H        | 01H         | RECYCLEQ.LD_BLOCK_ST_FORWARD       | Counts the number of occurrences a retired load gets blocked because its address partially overlaps with a store.                                  | PSDLA        |
| 03H        | 02H         | RECYCLEQ.LD_BLOCK_STD_NOT_READY    | Counts the number of occurrences a retired load gets blocked because its address overlaps with a store whose data is not ready.                    |              |
| 03H        | 04H         | RECYCLEQ.ST_SPLITS                 | Counts the number of occurrences a retired store that is a cache line split. Each split should be counted only once.                               |              |
| 03H        | 08H         | RECYCLEQ.LD_SPLITS                 | Counts the number of occurrences a retired load that is a cache line split. Each split should be counted only once.                                | PSDLA        |
| 03H        | 10H         | RECYCLEQ.LOCK                      | Counts all the retired locked loads. It does not include stores because we would double count if we count stores.                                  |              |
| 03H        | 20H         | RECYCLEQ.STA_FULL                  | Counts the store micro-ops retired that were pushed in the recycle queue because the store address buffer is full.                                 |              |
| 03H        | 40H         | RECYCLEQ.ANY_LD                    | Counts any retired load that was pushed into the recycle queue for any reason.   |              |
| 03H        | 80H         | RECYCLEQ.ANY_ST                    | Counts any retired store that was pushed into the recycle queue for any reason.  |              |
| 04H        | 01H         | MEM_UOPS_RETIRED.L1_MISS_LOADS     | Counts the number of load micro-ops retired that miss in L1 D cache.   |              |
| 04H        | 02H         | MEM_UOPS_RETIRED.L2_HIT_LOADS      | Counts the number of load micro-ops retired that hit in the L2.  | PSDLA        |
| 04H        | 04H         | MEM_UOPS_RETIRED.L2_MISS_LOADS     | Counts the number of load micro-ops retired that miss in the L2.   | PSDLA        |
| 04H        | 08H         | MEM_UOPS_RETIRED.DTLB_MISSES_LOADS | Counts the number of load micro-ops retired that cause a DTLB miss.  | PSDLA        |
| 04H        | 10H         | MEM_UOPS_RETIRED.UTLB_MISSES_LOADS | Counts the number of load micro-ops retired that caused micro TLB miss.  |              |
| 04H        | 20H         | MEM_UOPS_RETIRED.HITM              | Counts the loads retired that get the data from the other core in the same tile in M state.  |              |
| 04H        | 40H         | MEM_UOPS_RETIRED.ALL_LOADS         | Counts all the load micro-ops retired.   |              |
| 04H        | 80H         | MEM_UOPS_RETIRED.ALL_STORES        | Counts all the store micro-ops retired.  |              |
| 05H        | 01H         | PAGE_WALKS.D_SIDE_WALKS            | Counts the total D-side page walks that are completed or started. The page walks started in the speculative path will also be counted.             | EdgeDetect=1 |
| 05H        | 01H         | PAGE_WALKS.D_SIDE_CYCLES           | Counts the total number of core cycles for all the D-side page walks. The cycles for page walks started in speculative path will also be included. |              |
| 05H        | 02H         | PAGE_WALKS.I_SIDE_WALKS            | Counts the total I-side page walks that are completed.   | EdgeDetect=1 |

**Table 19-6. Performance Events of the Processor Core Supported by Knights Landing and Knights Mill Microarchitectures**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description   | Comment   |
|------------|-------------|--|---|---|
| 05H        | 02H         | PAGE_WALKS.I_SIDE_CYCLES               | Counts the total number of core cycles for all the I-side page walks. The cycles for page walks started in speculative path will also be included.  |   |
| 05H        | 03H         | PAGE_WALKS.WALKS                       | Counts the total page walks that are completed (I-side and D-side).   | EdgeDetect=1  |
| 05H        | 03H         | PAGE_WALKS.CYCLES                      | Counts the total number of core cycles for all the page walks. The cycles for page walks started in speculative path will also be included.   |   |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS                 | Counts the number of L2 cache misses. Also called L2_REQUESTS_MISS.   |   |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE            | Counts the total number of L2 cache references. Also called L2_REQUESTS_REFERENCE.  |   |
| 30H        | 00H         | L2_REQUESTS_REJECT.ALL                 | Counts the number of MEC requests from the L2Q that reference a cache line (cacheable requests) excluding SW prefetches filling only to L2 cache and L1 evictions (automatically excludes L2HWP, UC, WC) that were rejected - Multiple repeated rejects should be counted multiple times. |   |
| 31H        | 00H         | CORE_REJECT_L2Q.ALL                    | Counts the number of MEC requests that were not accepted into the L2Q because of any L2 queue reject condition. There is no concept of at-ret here. It might include requests due to instructions in the speculative path.  |   |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P              | Counts the number of unhalted core clock cycles.  |   |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF                   | Counts the number of unhalted reference clock cycles.   |   |
| 3EH        | 04H         | L2_PREFETCHER.ALLOC_XQ                 | Counts the number of L2HWP allocated into XQ GP.  |   |
| 80H        | 01H         | ICACHE.HIT                             | Counts all instruction fetches that hit the instruction cache.  |   |
| 80H        | 02H         | ICACHE.MISSES                          | Counts all instruction fetches that miss the instruction cache or produce memory requests. An instruction fetch miss is counted only once and not once for every cycle it is outstanding.   |   |
| 80H        | 03H         | ICACHE.ACCESSSES                       | Counts all instruction fetches, including uncacheable fetches.  |   |
| 86H        | 04H         | FETCH_STALL.ICACHE_FILL_PENDING_CYCLES | Counts the number of core cycles the fetch stalls because of an icache miss. This is a cumulative count of core cycles the fetch stalled for all icache misses.   |   |
| B7H        | 01H         | OFFCORE_RESPONSE_0                     | See Section 18.4.1.1.2.   | Requires MSR_OFFCORE_RESP 0 to specify request type and response. |
| B7H        | 02H         | OFFCORE_RESPONSE_1                     | See Section 18.4.1.1.2.   | Requires MSR_OFFCORE_RESP 1 to specify request type and response. |
| COH        | 00H         | INST_RETIRED.ANY_P                     | Counts the total number of instructions retired.  | PS  |

**Table 19-6. Performance Events of the Processor Core Supported by Knights Landing and Knights Mill Microarchitectures**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment |
|------------|-------------|--------------------------------|--|---------|
| C2H        | 01H         | UOPS_RETIRED.MS                | Counts the number of micro-ops retired that are from the complex flows issued by the micro-sequencer (MS).   |         |
| C2H        | 10H         | UOPS_RETIRED.ALL               | Counts the number of micro-ops retired.  |         |
| C2H        | 20H         | UOPS_RETIRED.SCALAR_SIMD       | Counts the number of scalar SSE, AVX, AVX2, and AVX-512 micro-ops except for loads (memory-to-register mov-type micro ops), division and sqrt.   |         |
| C2H        | 40H         | UOPS_RETIRED.PACKED_SIMD       | Counts the number of packed SSE, AVX, AVX2, and AVX-512 micro-ops (both floating point and integer) except for loads (memory-to-register mov-type micro-ops), packed byte and word multiplies. |         |
| C3H        | 01H         | MACHINE_CLEARS.SMC             | Counts the number of times that the machine clears due to program modifying data within 1K of a recently fetched code page.  |         |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of times the machine clears due to memory ordering hazards.  |         |
| C3H        | 04H         | MACHINE_CLEARS.FP_ASSIST       | Counts the number of floating operations retired that required microcode assists.  |         |
| C3H        | 08H         | MACHINE_CLEARS.ALL             | Counts all machine clears.   |         |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts the number of branch instructions retired.  | PS      |
| C4H        | 7EH         | BR_INST_RETIRED.JCC            | Counts the number of JCC branch instructions retired.  | PS      |
| C4H        | BFH         | BR_INST_RETIRED.FAR_BRANCH     | Counts the number of far branch instructions retired.  | PS      |
| C4H        | EBH         | BR_INST_RETIRED.NON_RETURN_IND | Counts the number of branch instructions retired that were near indirect CALL or near indirect JMP.  | PS      |
| C4H        | F7H         | BR_INST_RETIRED.RETURN         | Counts the number of near RET branch instructions retired.   | PS      |
| C4H        | F9H         | BR_INST_RETIRED.CALL           | Counts the number of near CALL branch instructions retired.  | PS      |
| C4H        | FBH         | BR_INST_RETIRED.IND_CALL       | Counts the number of near indirect CALL branch instructions retired.   | PS      |
| C4H        | FDH         | BR_INST_RETIRED.REL_CALL       | Counts the number of near relative CALL branch instructions retired.   | PS      |
| C4H        | FEH         | BR_INST_RETIRED.TAKEN_JCC      | Counts the number of branch instructions retired that were taken conditional jumps.  | PS      |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES   | Counts the number of mispredicted branch instructions retired.   | PS      |
| C5H        | 7EH         | BR_MISP_RETIRED.JCC            | Counts the number of mispredicted JCC branch instructions retired.   | PS      |
| C5H        | BFH         | BR_MISP_RETIRED.FAR_BRANCH     | Counts the number of mispredicted far branch instructions retired.   | PS      |
| C5H        | EBH         | BR_MISP_RETIRED.NON_RETURN_IND | Counts the number of mispredicted branch instructions retired that were near indirect CALL or near indirect JMP.   | PS      |
| C5H        | F7H         | BR_MISP_RETIRED.RETURN         | Counts the number of mispredicted near RET branch instructions retired.  | PS      |



**Table 19-6. Performance Events of the Processor Core Supported by Knights Landing and Knights Mill Microarchitectures**

| Event Num.   | Umask Value | Event Mask Mnemonic           | Description   | Comment |
|--|-------------|-------------------------------|---|---------|
| C5H  | F9H         | BR_MISP_RETIREDCALL           | Counts the number of mispredicted near CALL branch instructions retired.  | PS      |
| C5H  | FBH         | BR_MISP_RETIREDIND_CALL       | Counts the number of mispredicted near indirect CALL branch instructions retired.   | PS      |
| C5H  | FDH         | BR_MISP_RETIREDREL_CALL       | Counts the number of mispredicted near relative CALL branch instructions retired.   | PS      |
| C5H  | FEH         | BR_MISP_RETIREDTAKEN_JCC      | Counts the number of mispredicted branch instructions retired that were taken conditional jumps.  | PS      |
| CAH  | 01H         | NO_ALLOC_CYCLES.ROB_FULL      | Counts the number of core cycles when no micro-ops are allocated and the ROB is full.   |         |
| CAH  | 04H         | NO_ALLOC_CYCLES.MISPREDICTS   | Counts the number of core cycles when no micro-ops are allocated and the alloc pipe is stalled waiting for a mispredicted branch to retire.   |         |
| CAH  | 20H         | NO_ALLOC_CYCLES.RAT_STALL     | Counts the number of core cycles when no micro-ops are allocated and a RATstall (caused by reservation station full) is asserted.             |         |
| CAH  | 90H         | NO_ALLOC_CYCLES.NOT_DELIVERED | Counts the number of core cycles when no micro-ops are allocated, the IQ is empty, and no other condition is blocking allocation.             |         |
| CAH  | 7FH         | NO_ALLOC_CYCLES.ALL           | Counts the total number of core cycles when no micro-ops are allocated for any reason.  |         |
| CBH  | 01H         | RS_FULL_STALL.MEC             | Counts the number of core cycles when allocation pipeline is stalled and is waiting for a free MEC reservation station entry.                 |         |
| CBH  | 1FH         | RS_FULL_STALL.ALL             | Counts the total number of core cycles the allocation pipeline is stalled when any one of the reservation stations is full.                   |         |
| CDH  | 01H         | CYCLES_DIV_BUSY.ALL           | Cycles the number of core cycles when divider is busy. Does not imply a stall waiting for the divider.  |         |
| E6H  | 01H         | BACLEARS.ALL                  | Counts the number of times the front end resteers for any branch as a result of another branch handling mechanism in the front end.           |         |
| E6H  | 08H         | BACLEARS.RETURN               | Counts the number of times the front end resteers for RET branches as a result of another branch handling mechanism in the front end.         |         |
| E6H  | 10H         | BACLEARS.COND                 | Counts the number of times the front end resteers for conditional branches as a result of another branch handling mechanism in the front end. |         |
| E7H  | 01H         | MS_DECODED.MS_ENTRY           | Counts the number of times the MSROM starts a flow of uops.   |         |
| PS: Also supports PEBS.<br>PSDLA: Also supports PEBS and DataLA. |             |                               |   |         |

## 19.5 PERFORMANCE MONITORING EVENTS FOR THE INTEL® CORE™ M AND 5TH GENERATION INTEL® CORE™ PROCESSORS

The Intel® Core™ M processors, the 5th generation Intel® Core™ processors and the Intel Xeon processor E3 1200 v4 product family are based on the Broadwell microarchitecture. They support the architectural performance monitoring events listed in Table 19-1. Model-specific performance monitoring events in the processor core are listed in Table 19-7. The events in Table 19-7 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_3DH and 06\_47H. Table 19-10 lists performance events supporting Intel TSX (see Section 18.3.6.5) and the events are available on processors based on Broadwell microarchitecture. Fixed counters in the core PMU support the architecture events defined in Table 19-2.

Model-specific performance monitoring events that are located in the uncore sub-system are implementation specific between different platforms using processors based on Broadwell microarchitecture and with different DisplayFamily\_DisplayModel signatures. Processors with CPUID signature of DisplayFamily\_DisplayModel 06\_3DH and 06\_47H support uncore performance events listed in Table 19-11.

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description   | Comment   |
|------------|-------------|-------------------------------------|---|---|
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD             | Loads blocked by overlapping with store buffer that cannot be forwarded.  |   |
| 03H        | 08H         | LD_BLOCKS.NO_SR                     | The number of times that split load operations are temporarily blocked because all resources for handling the split accesses are in use.                                  |   |
| 05H        | 01H         | MISALIGN_MEM_REF.LOADS              | Speculative cache-line split load uops dispatched to L1D.   |   |
| 05H        | 02H         | MISALIGN_MEM_REF.STORES             | Speculative cache-line split store-address uops dispatched to L1D.  |   |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS     | False dependencies in MOB due to partial compare on address.  |   |
| 08H        | 01H         | DTLB_LOAD_MISSES.MISS_CAUSES_A_WALK | Load misses in all TLB levels that cause a page walk of any page size.  |   |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED_4K  | Completed page walks due to demand load misses that caused 4K page walks in any TLB levels.   |   |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_DURATION      | Cycle PMH is busy with a walk.  |   |
| 08H        | 20H         | DTLB_LOAD_MISSES.STLB_HIT_4K        | Load misses that missed DTLB but hit STLB (4K).   |   |
| 0DH        | 03H         | INT_MISC.RECOVERY_CYCLES            | Cycles waiting to recover after Machine Clears except JEClear. Set Cmask= 1.  | Set Edge to count occurrences.                  |
| 0EH        | 01H         | UOPS_ISSUED.ANY                     | Increments each cycle the # of uops issued by the RAT to RS. Set Cmask = 1, Inv = 1, Any= 1 to count stalled cycles of this core.   | Set Cmask = 1, Inv = 1 to count stalled cycles. |
| 0EH        | 10H         | UOPS_ISSUED.FLAGS_MERGE             | Number of flags-merge uops allocated. Such uops add delay.  |   |
| 0EH        | 20H         | UOPS_ISSUED.SLOW_LEA                | Number of slow LEA or similar uops allocated. Such uop has 3 sources (for example, 2 sources + immediate) regardless of whether it is a result of LEA instruction or not. |   |
| 0EH        | 40H         | UOPS_ISSUED.SINGLE_MUL              | Number of multiply packed/scalar single precision uops allocated.   |   |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description  | Comment   |
|------------|-------------|--------------------------------------|--|---|
| 14H        | 01H         | ARITH.FPU_DIV_ACTIVE                 | Cycles when divider is busy executing divide operations.   |   |
| 24H        | 21H         | L2_RQSTS.DEMAND_DATA_RD_MISS         | Demand data read requests that missed L2, no rejects.  |   |
| 24H        | 41H         | L2_RQSTS.DEMAND_DATA_RD_HIT          | Demand data read requests that hit L2 cache.   |   |
| 24H        | 50H         | L2_RQSTS.L2_PF_HIT                   | Counts all L2 HW prefetcher requests that hit L2.  |   |
| 24H        | 30H         | L2_RQSTS.L2_PF_MISS                  | Counts all L2 HW prefetcher requests that missed L2.   |   |
| 24H        | E1H         | L2_RQSTS.ALL_DEMAND_DATA_RD          | Counts any demand and L1 HW prefetch data load requests to L2.   |   |
| 24H        | E2H         | L2_RQSTS.ALL_RFO                     | Counts all L2 store RFO requests.  |   |
| 24H        | E4H         | L2_RQSTS.ALL_CODE_RD                 | Counts all L2 code requests.   |   |
| 24H        | F8H         | L2_RQSTS.ALL_PF                      | Counts all L2 HW prefetcher requests.  |   |
| 27H        | 50H         | L2_DEMAND_RQSTS.WB_HIT               | Not rejected writebacks that hit L2 cache.   |   |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE          | This event counts requests originating from the core that reference a cache line in the last level cache.  | See Table 19-1.                                   |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS               | This event counts each cache miss condition for references to the last level cache.  | See Table 19-1.                                   |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P            | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. | See Table 19-1.                                   |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK     | Increments at the frequency of XCLK (100 MHz) when not halted.   | See Table 19-1.                                   |
| 48H        | 01H         | L1D_PEND_MISS.PENDING                | Increments the number of outstanding L1D misses every cycle. Set Cmask = 1 and Edge = 1 to count occurrences.  | Counter 2 only.<br>Set Cmask = 1 to count cycles. |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK | Miss in all TLB levels causes a page walk of any page size (4K/2M/4M/1G).  |   |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED_4K  | Completed page walks due to store misses in one or more TLB levels of 4K page structure.   |   |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_DURATION      | Cycles PMH is busy with this walk.   |   |
| 49H        | 20H         | DTLB_STORE_MISSES.STLB_HIT_4K        | Store misses that missed DTLB but hit STLB (4K).   |   |
| 4CH        | 02H         | LOAD_HIT_PRE.HW_PF                   | Non-SW-prefetch load dispatches that hit fill buffer allocated for H/W prefetch.   |   |
| 4FH        | 10H         | EPT.WALK_CYCLES                      | Cycles of Extended Page Table walks.   |   |
| 51H        | 01H         | L1D.REPLACEMENT                      | Counts the number of lines brought into the L1 data cache.   |   |
| 58H        | 04H         | MOVE_ELIMINATION.INT_NOT_ELIMINATED  | Number of integer move elimination candidate uops that were not eliminated.  |   |
| 58H        | 08H         | MOVE_ELIMINATION.SIMD_NOT_ELIMINATED | Number of SIMD move elimination candidate uops that were not eliminated.   |   |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description  | Comment                        |
|------------|-------------|---|--|--------------------------------|
| 58H        | 01H         | MOVE_ELIMINATION.INT_ELIMINATED             | Number of integer move elimination candidate uops that were eliminated.  |                                |
| 58H        | 02H         | MOVE_ELIMINATION.SIMD_ELIMINATED            | Number of SIMD move elimination candidate uops that were eliminated.   |                                |
| 5CH        | 01H         | CPL_CYCLES.RING0                            | Unhalted core cycles when the thread is in ring 0.   | Use Edge to count transition.  |
| 5CH        | 02H         | CPL_CYCLES.RING123                          | Unhalted core cycles when the thread is not in ring 0.   |                                |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES                      | Cycles the RS is empty for the thread.   |                                |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD | Offcore outstanding demand data read transactions in SQ to uncore. Set Cmask=1 to count cycles.  | Use only when HTT is off.      |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_CODE_RD | Offcore outstanding demand code read transactions in SQ to uncore. Set Cmask=1 to count cycles.  | Use only when HTT is off.      |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO     | Offcore outstanding RFO store transactions in SQ to uncore. Set Cmask=1 to count cycles.   | Use only when HTT is off.      |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD    | Offcore outstanding cacheable data read transactions in SQ to uncore. Set Cmask=1 to count cycles.                                     | Use only when HTT is off.      |
| 63H        | 01H         | LOCK_CYCLES.SPLIT_LOCK_UC_LOCK_DURATION     | Cycles in which the L1D and L2 are locked, due to a UC lock or split lock.   |                                |
| 63H        | 02H         | LOCK_CYCLES.CACHE_LOCK_DURATION             | Cycles in which the L1D is locked.   |                                |
| 79H        | 02H         | IDQ.EMPTY                                   | Counts cycles the IDQ is empty.  |                                |
| 79H        | 04H         | IDQ.MITE_UOPS                               | Increment each cycle # of uops delivered to IDQ from MITE path. Set Cmask = 1 to count cycles.   | Can combine Umask 04H and 20H. |
| 79H        | 08H         | IDQ.DSB_UOPS                                | Increment each cycle # of uops delivered to IDQ from DSB path. Set Cmask = 1 to count cycles.  | Can combine Umask 08H and 10H. |
| 79H        | 10H         | IDQ.MS_DSB_UOPS                             | Increment each cycle # of uops delivered to IDQ when MS_busy by DSB. Set Cmask = 1 to count cycles. Add Edge=1 to count # of delivery. | Can combine Umask 04H, 08H.    |
| 79H        | 20H         | IDQ.MS_MITE_UOPS                            | Increment each cycle # of uops delivered to IDQ when MS_busy by MITE. Set Cmask = 1 to count cycles.                                   | Can combine Umask 04H, 08H.    |
| 79H        | 30H         | IDQ.MS_UOPS                                 | Increment each cycle # of uops delivered to IDQ from MS by either DSB or MITE. Set Cmask = 1 to count cycles.                          | Can combine Umask 04H, 08H.    |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_ANY_UOPS                 | Counts cycles DSB is delivered at least one uops. Set Cmask = 1.   |                                |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_4_UOPS                   | Counts cycles DSB is delivered four uops. Set Cmask = 4.   |                                |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_ANY_UOPS                | Counts cycles MITE is delivered at least one uop. Set Cmask = 1.   |                                |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_4_UOPS                  | Counts cycles MITE is delivered four uops. Set Cmask = 4.  |                                |
| 79H        | 3CH         | IDQ.MITE_ALL_UOPS                           | Number of uops delivered to IDQ from any path.   |                                |
| 80H        | 02H         | ICACHE.MISSES                               | Number of Instruction Cache, Streaming Buffer and Victim Cache Misses. Includes UC accesses.   |                                |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description   | Comment                           |
|------------|-------------|--|---|-----------------------------------|
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK         | Misses in ITLB that cause a page walk of any page size.   |                                   |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETE_D_4K         | Completed page walks due to misses in ITLB 4K page entries.   |                                   |
| 85H        | 10H         | ITLB_MISSES.WALK_DURATION              | Cycle PMH is busy with a walk.  |                                   |
| 85H        | 20H         | ITLB_MISSES.STLB_HIT_4K                | ITLB misses that hit STLB (4K).   |                                   |
| 87H        | 01H         | ILD_STALL.LCP                          | Stalls caused by changing prefix length of the instruction.   |                                   |
| 88H        | 01H         | BR_INST_EXEC.COND                      | Qualify conditional near branch instructions executed, but not necessarily retired.                                       | Must combine with umask 40H, 80H. |
| 88H        | 02H         | BR_INST_EXEC.DIRECT_JMP                | Qualify all unconditional near branch instructions excluding calls and indirect branches.                                 | Must combine with umask 80H.      |
| 88H        | 04H         | BR_INST_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify executed indirect near branch instructions that are not calls or returns.   | Must combine with umask 80H.      |
| 88H        | 08H         | BR_INST_EXEC.RETURN_NEAR               | Qualify indirect near branches that have a return mnemonic.   | Must combine with umask 80H.      |
| 88H        | 10H         | BR_INST_EXEC.DIRECT_NEAR_CALL          | Qualify unconditional near call branch instructions, excluding non-call branch, executed.                                 | Must combine with umask 80H.      |
| 88H        | 20H         | BR_INST_EXEC.INDIRECT_NEAR_CALL        | Qualify indirect near calls, including both register and memory indirect, executed.                                       | Must combine with umask 80H.      |
| 88H        | 40H         | BR_INST_EXEC.NONTAKEN                  | Qualify non-taken near branches executed.   | Applicable to umask 01H only.     |
| 88H        | 80H         | BR_INST_EXEC.TAKEN                     | Qualify taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                                      |                                   |
| 88H        | FFH         | BR_INST_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 89H        | 01H         | BR_MISP_EXEC.COND                      | Qualify conditional near branch instructions mispredicted.  | Must combine with umask 40H, 80H. |
| 89H        | 04H         | BR_MISP_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify mispredicted indirect near branch instructions that are not calls or returns.                                     | Must combine with umask 80H.      |
| 89H        | 08H         | BR_MISP_EXEC.RETURN_NEAR               | Qualify mispredicted indirect near branches that have a return mnemonic.  | Must combine with umask 80H.      |
| 89H        | 10H         | BR_MISP_EXEC.DIRECT_NEAR_CALL          | Qualify mispredicted unconditional near call branch instructions, excluding non-call branch, executed.                    | Must combine with umask 80H.      |
| 89H        | 20H         | BR_MISP_EXEC.INDIRECT_NEAR_CALL        | Qualify mispredicted indirect near calls, including both register and memory indirect, executed.                          | Must combine with umask 80H.      |
| 89H        | 40H         | BR_MISP_EXEC.NONTAKEN                  | Qualify mispredicted non-taken near branches executed.  | Applicable to umask 01H only.     |
| 89H        | 80H         | BR_MISP_EXEC.TAKEN                     | Qualify mispredicted taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                         |                                   |
| 89H        | FFH         | BR_MISP_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.COORE           | Count issue pipeline slots where no uop was delivered from the front end to the back end when there is no back end stall. | Use Cmask to qualify uop b/w.     |
| A1H        | 01H         | UOPS_DISPATCHED_PORT.PORT_0            | Counts the number of cycles in which a uop is dispatched to port 0.   | Set AnyThread to count per core.  |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description  | Comment                          |
|------------|-------------|----------------------------------|--|----------------------------------|
| A1H        | 02H         | UOPS_DISPATCHED_PORT.PORT_1      | Counts the number of cycles in which a uop is dispatched to port 1.                  | Set AnyThread to count per core. |
| A1H        | 04H         | UOPS_DISPATCHED_PORT.PORT_2      | Counts the number of cycles in which a uop is dispatched to port 2.                  | Set AnyThread to count per core. |
| A1H        | 08H         | UOPS_DISPATCHED_PORT.PORT_3      | Counts the number of cycles in which a uop is dispatched to port 3.                  | Set AnyThread to count per core. |
| A1H        | 10H         | UOPS_DISPATCHED_PORT.PORT_4      | Counts the number of cycles in which a uop is dispatched to port 4.                  | Set AnyThread to count per core. |
| A1H        | 20H         | UOPS_DISPATCHED_PORT.PORT_5      | Counts the number of cycles in which a uop is dispatched to port 5.                  | Set AnyThread to count per core. |
| A1H        | 40H         | UOPS_DISPATCHED_PORT.PORT_6      | Counts the number of cycles in which a uop is dispatched to port 6.                  | Set AnyThread to count per core. |
| A1H        | 80H         | UOPS_DISPATCHED_PORT.PORT_7      | Counts the number of cycles in which a uop is dispatched to port 7.                  | Set AnyThread to count per core. |
| A2H        | 01H         | RESOURCE_STALLS.ANY              | Cycles Allocation is stalled due to resource related reason.                         |                                  |
| A2H        | 04H         | RESOURCE_STALLS.RS               | Cycles stalled due to no eligible RS entry available.                                |                                  |
| A2H        | 08H         | RESOURCE_STALLS.SB               | Cycles stalled due to no store buffers available (not including draining form sync). |                                  |
| A2H        | 10H         | RESOURCE_STALLS.ROB              | Cycles stalled due to re-order buffer full.  |                                  |
| A8H        | 01H         | LSD.UOPS                         | Number of uops delivered by the LSD.   |                                  |
| ABH        | 02H         | DSB2MITE_SWITCHES.PENALTY_CYCLES | Cycles of delay due to Decode Stream Buffer to MITE switches.                        |                                  |
| AEH        | 01H         | ITLB.ITLB_FLUSH                  | Counts the number of ITLB flushes; includes 4k/2M/4M pages.                          |                                  |
| B0H        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD  | Demand data read requests sent to uncore.  | Use only when HTT is off.        |
| B0H        | 02H         | OFFCORE_REQUESTS.DEMAND_CODE_RD  | Demand code read requests sent to uncore.  | Use only when HTT is off.        |
| B0H        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO      | Demand RFO read requests sent to uncore, including regular RFOs, locks, ItoM.        | Use only when HTT is off.        |
| B0H        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD     | Data read requests sent to uncore (demand and prefetch).                             | Use only when HTT is off.        |
| B1H        | 01H         | UOPS_EXECUTED.THREAD             | Counts total number of uops to be executed per-logical-processor each cycle.         | Use Cmask to count stall cycles. |
| B1H        | 02H         | UOPS_EXECUTED.CORE               | Counts total number of uops to be executed per-core each cycle.                      | Do not need to set ANY.          |
| B7H        | 01H         | OFF_CORE_RESPONSE_0              | See Section 18.3.4.5, "Off-core Response Performance Monitoring".                    | Requires MSR 01A6H.              |
| BBH        | 01H         | OFF_CORE_RESPONSE_1              | See Section 18.3.4.5, "Off-core Response Performance Monitoring".                    | Requires MSR 01A7H.              |
| BCH        | 11H         | PAGE_WALKER_LOADS.DTLB_L1        | Number of DTLB page walker loads that hit in the L1+FB.                              |                                  |
| BCH        | 21H         | PAGE_WALKER_LOADS.ITLB_L1        | Number of ITLB page walker loads that hit in the L1+FB.                              |                                  |
| BCH        | 12H         | PAGE_WALKER_LOADS.DTLB_L2        | Number of DTLB page walker loads that hit in the L2.                                 |                                  |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment  |
|------------|-------------|--------------------------------|--|--|
| BCH        | 22H         | PAGE_WALKER_LOADS.ITLB_L2      | Number of ITLB page walker loads that hit in the L2.   |  |
| BCH        | 14H         | PAGE_WALKER_LOADS.DTLB_L3      | Number of DTLB page walker loads that hit in the L3.   |  |
| BCH        | 24H         | PAGE_WALKER_LOADS.ITLB_L3      | Number of ITLB page walker loads that hit in the L3.   |  |
| BCH        | 18H         | PAGE_WALKER_LOADS.DTLB_MEMORY  | Number of DTLB page walker loads from memory.  |  |
| COH        | 00H         | INST_RETIRED.ANY_P             | Number of instructions at retirement.  | See Table 19-1.  |
| COH        | 01H         | INST_RETIRED.PREC_DIST         | Precise instruction retired event with HW to reduce effect of PEBS shadow in IP distribution.                                | PMC1 only.   |
| COH        | 02H         | INST_RETIRED.X87               | FP operations retired. X87 FP operations that have no exceptions.  |  |
| C1H        | 08H         | OTHER_ASSISTS.AVX_TO_SSE       | Number of transitions from AVX-256 to legacy SSE when penalty applicable.  |  |
| C1H        | 10H         | OTHER_ASSISTS.SSE_TO_AVX       | Number of transitions from SSE to AVX-256 when penalty applicable.   |  |
| C1H        | 40H         | OTHER_ASSISTS.ANY_WB_ASSIST    | Number of microcode assists invoked by HW upon uop writeback.  |  |
| C2H        | 01H         | UOPS_RETIRED.ALL               | Counts the number of micro-ops retired.<br>Use cmask=1 and invert to count active cycles or stalled cycles.                  | Supports PEBS and DataLA, use Any=1 for core granular. |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS      | Counts the number of retirement slots used each cycle.   | Supports PEBS.   |
| C3H        | 01H         | MACHINE_CLEARS.CYCLES          | Counts cycles while a machine clears stalled forward progress of a logical processor or a processor core.                    |  |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of machine clears due to memory order conflicts.   |  |
| C3H        | 04H         | MACHINE_CLEARS.SMC             | Number of self-modifying-code machine clears detected.   |  |
| C3H        | 20H         | MACHINE_CLEARS.MASKMOV         | Counts the number of executed AVX masked load operations that refer to an illegal address range with the mask bits set to 0. |  |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Branch instructions at retirement.   | See Table 19-1.  |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL    | Counts the number of conditional branch instructions retired.  | Supports PEBS.   |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL      | Direct and indirect near call instructions retired.  | Supports PEBS.   |
| C4H        | 04H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts the number of branch instructions retired.  | Supports PEBS.   |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETURN    | Counts the number of near return instructions retired.   | Supports PEBS.   |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN      | Counts the number of not taken branch instructions retired.  |  |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKEN     | Number of near taken branches retired.   | Supports PEBS.   |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANCHES   | Number of far branches retired.  |  |

**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description   | Comment                        |
|------------|-------------|--------------------------------------|---|--------------------------------|
| C5H        | 00H         | BR_MISP_RETIRE.ALL_BRANC<br>HES      | Mispredicted branch instructions at retirement.   | See Table 19-1.                |
| C5H        | 01H         | BR_MISP_RETIRE.CONDITIONA<br>L       | Mispredicted conditional branch instructions retired.   | Supports PEBS.                 |
| C5H        | 04H         | BR_MISP_RETIRE.ALL_BRANC<br>HES      | Mispredicted macro branch instructions retired.   | Supports PEBS.                 |
| CAH        | 02H         | FP_ASSIST.X87_OUTPUT                 | Number of X87 FP assists due to output values.  |                                |
| CAH        | 04H         | FP_ASSIST.X87_INPUT                  | Number of X87 FP assists due to input values.   |                                |
| CAH        | 08H         | FP_ASSIST.SIMD_OUTPUT                | Number of SIMD FP assists due to output values.   |                                |
| CAH        | 10H         | FP_ASSIST.SIMD_INPUT                 | Number of SIMD FP assists due to input values.  |                                |
| CAH        | 1EH         | FP_ASSIST.ANY                        | Cycles with any input/output SSE* or FP assists.  |                                |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSERTS          | Count cases of saving new LBR records by hardware.  |                                |
| CDH        | 01H         | MEM_TRANS_RETIRE.LOAD_L<br>ATENCY    | Randomly sampled loads whose latency is above a user defined threshold. A small fraction of the overall loads are sampled due to randomization. | Specify threshold in MSR 3F6H. |
| DOH        | 11H         | MEM_UOPS_RETIRE.STLB_MIS<br>S_LOADS  | Retired load uops that miss the STLB.   | Supports PEBS and DataLA.      |
| DOH        | 12H         | MEM_UOPS_RETIRE.STLB_MIS<br>S_STORES | Retired store uops that miss the STLB.  | Supports PEBS and DataLA.      |
| DOH        | 21H         | MEM_UOPS_RETIRE.LOCK_LOADS           | Retired load uops with locked access.   | Supports PEBS and DataLA.      |
| DOH        | 41H         | MEM_UOPS_RETIRE.SPLIT_LOADS          | Retired load uops that split across a cacheline boundary.   | Supports PEBS and DataLA.      |
| DOH        | 42H         | MEM_UOPS_RETIRE.SPLIT_STORES         | Retired store uops that split across a cacheline boundary.  | Supports PEBS and DataLA.      |
| DOH        | 81H         | MEM_UOPS_RETIRE.ALL_LOADS            | All retired load uops.  | Supports PEBS and DataLA.      |
| DOH        | 82H         | MEM_UOPS_RETIRE.ALL_STORES           | All retired store uops.   | Supports PEBS and DataLA.      |
| D1H        | 01H         | MEM_LOAD_UOPS_RETIRE.L1_HIT          | Retired load uops with L1 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 02H         | MEM_LOAD_UOPS_RETIRE.L2_HIT          | Retired load uops with L2 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 04H         | MEM_LOAD_UOPS_RETIRE.L3_HIT          | Retired load uops with L3 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 08H         | MEM_LOAD_UOPS_RETIRE.L1_MISS         | Retired load uops missed L1 cache as data sources.  | Supports PEBS and DataLA.      |
| D1H        | 10H         | MEM_LOAD_UOPS_RETIRE.L2_MISS         | Retired load uops missed L2. Unknown data source excluded.  | Supports PEBS and DataLA.      |
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRE.L3_MISS         | Retired load uops missed L3. Excludes unknown data source.  | Supports PEBS and DataLA.      |
| D1H        | 40H         | MEM_LOAD_UOPS_RETIRE.HIT_LFB         | Retired load uops where data sources were load uops missed L1 but hit FB due to preceding miss to the same cache line with data not ready.      | Supports PEBS and DataLA.      |



**Table 19-7. Performance Events of the Processor Core Supported by Broadwell Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                          | Description  | Comment                          |
|------------|-------------|--|--|----------------------------------|
| D2H        | 01H         | MEM_LOAD_UOPS_L3_HIT_RETI<br>RED.XSNP_MISS   | Retired load uops where data sources were L3 hit and cross-core snoop missed in on-pkg core cache. | Supports PEBS and DataLA.        |
| D2H        | 02H         | MEM_LOAD_UOPS_L3_HIT_RETI<br>RED.XSNP_HIT    | Retired load uops where data sources were L3 and cross-core snoop hits in on-pkg core cache.       | Supports PEBS and DataLA.        |
| D2H        | 04H         | MEM_LOAD_UOPS_L3_HIT_RETI<br>RED.XSNP_HITM   | Retired load uops where data sources were HitM responses from shared L3.                           | Supports PEBS and DataLA.        |
| D2H        | 08H         | MEM_LOAD_UOPS_L3_HIT_RETI<br>RED.XSNP_NONE   | Retired load uops where data sources were hits in L3 without snoops required.                      | Supports PEBS and DataLA.        |
| D3H        | 01H         | MEM_LOAD_UOPS_L3_MISS_RE<br>TIRED.LOCAL_DRAM | Retired load uops where data sources missed L3 but serviced from local dram.                       | Supports PEBS and DataLA.        |
| FOH        | 01H         | L2_TRANS.DEMAND_DATA_RD                      | Demand data read requests that access L2 cache.  |                                  |
| FOH        | 02H         | L2_TRANS.RFO                                 | RFO requests that access L2 cache.   |                                  |
| FOH        | 04H         | L2_TRANS.CODE_RD                             | L2 cache accesses when fetching instructions.  |                                  |
| FOH        | 08H         | L2_TRANS.ALL_PF                              | Any MLC or L3 HW prefetch accessing L2, including rejects.   |                                  |
| FOH        | 10H         | L2_TRANS.L1D_WB                              | L1D writebacks that access L2 cache.   |                                  |
| FOH        | 20H         | L2_TRANS.L2_FILL                             | L2 fill requests that access L2 cache.   |                                  |
| FOH        | 40H         | L2_TRANS.L2_WB                               | L2 writebacks that access L2 cache.  |                                  |
| FOH        | 80H         | L2_TRANS.ALL_REQUESTS                        | Transactions accessing L2 pipe.  |                                  |
| F1H        | 01H         | L2_LINES_IN.I                                | L2 cache lines in I state filling L2.  | Counting does not cover rejects. |
| F1H        | 02H         | L2_LINES_IN.S                                | L2 cache lines in S state filling L2.  | Counting does not cover rejects. |
| F1H        | 04H         | L2_LINES_IN.E                                | L2 cache lines in E state filling L2.  | Counting does not cover rejects. |
| F1H        | 07H         | L2_LINES_IN.ALL                              | L2 cache lines filling L2.   | Counting does not cover rejects. |
| F2H        | 05H         | L2_LINES_OUT.DEMAND_CLEAN                    | Clean L2 cache lines evicted by demand.  |                                  |

Table 19-10 lists performance events supporting Intel TSX (see Section 18.3.6.5) and the events are applicable to processors based on Broadwell microarchitecture. Where Broadwell microarchitecture implements TSX-related event semantics that differ from Table 19-10, they are listed in Table 19-8.

**Table 19-8. Intel® TSX Performance Event Addendum in Processors Based on Broadwell Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description   | Comment |
|------------|-------------|-----------------------|---|---------|
| 54H        | 02H         | TX_MEM.ABORT_CAPACITY | Number of times a transactional abort was signaled due to a data capacity limitation for transactional reads or writes. |         |

## 19.6 PERFORMANCE MONITORING EVENTS FOR THE 4TH GENERATION INTEL® CORE™ PROCESSORS

4th generation Intel® Core™ processors and Intel Xeon processor E3-1200 v3 product family are based on the Haswell microarchitecture. They support the architectural performance monitoring events listed in Table 19-1. Model-specific performance monitoring events in the processor core are listed in Table 19-9. The events in Table

19-9 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_3CH, 06\_45H and 06\_46H. Table 19-10 lists performance events focused on supporting Intel TSX (see Section 18.3.6.5). Fixed counters in the core PMU support the architecture events defined in Table 19-2.

Additional information on event specifics (e.g., derivative events using specific IA32\_PERFVTSELx modifiers, limitations, special notes and recommendations) can be found at <https://software.intel.com/en-us/forums/software-tuning-performance-optimization-platform-monitoring>.

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors**

| Event Num. | Umask Value | Event Mask Mnemonic                   | Description   | Comment   |
|------------|-------------|---------------------------------------|---|---|
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD               | Loads blocked by overlapping with store buffer that cannot be forwarded.  |   |
| 03H        | 08H         | LD_BLOCKS.NO_SR                       | The number of times that split load operations are temporarily blocked because all resources for handling the split accesses are in use.                                  |   |
| 05H        | 01H         | MISALIGN_MEM_REF.LOADS                | Speculative cache-line split load uops dispatched to L1D.   |   |
| 05H        | 02H         | MISALIGN_MEM_REF.STORES               | Speculative cache-line split store-address uops dispatched to L1D.  |   |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS       | False dependencies in MOB due to partial compare on address.  |   |
| 08H        | 01H         | DTLB_LOAD_MISSES.MISS_CAUSES_A_WALK   | Misses in all TLB levels that cause a page walk of any page size.   |   |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED_4K    | Completed page walks due to demand load misses that caused 4K page walks in any TLB levels.   |   |
| 08H        | 04H         | DTLB_LOAD_MISSES.WALK_COMPLETED_2M_4M | Completed page walks due to demand load misses that caused 2M/4M page walks in any TLB levels.  |   |
| 08H        | 0EH         | DTLB_LOAD_MISSES.WALK_COMPLETED       | Completed page walks in any TLB of any page size due to demand load misses.   |   |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_DURATION        | Cycle PMH is busy with a walk.  |   |
| 08H        | 20H         | DTLB_LOAD_MISSES.STLB_HIT_4K          | Load misses that missed DTLB but hit STLB (4K).   |   |
| 08H        | 40H         | DTLB_LOAD_MISSES.STLB_HIT_2M          | Load misses that missed DTLB but hit STLB (2M).   |   |
| 08H        | 60H         | DTLB_LOAD_MISSES.STLB_HIT             | Number of cache load STLB hits. No page walk.   |   |
| 08H        | 80H         | DTLB_LOAD_MISSES.PDE_CACHE_MISS       | DTLB demand load misses with low part of linear-to-physical address translation missed.   |   |
| 0DH        | 03H         | INT_MISC.RECOVERY_CYCLES              | Cycles waiting to recover after Machine Clears except JEClear. Set Cmask= 1.  | Set Edge to count occurrences.                  |
| 0EH        | 01H         | UOPS_ISSUED.ANY                       | Increments each cycle the # of uops issued by the RAT to RS. Set Cmask = 1, Inv = 1, Any= 1 to count stalled cycles of this core.   | Set Cmask = 1, Inv = 1 to count stalled cycles. |
| 0EH        | 10H         | UOPS_ISSUED.FLAGS_MERGE               | Number of flags-merge uops allocated. Such uops add delay.  |   |
| 0EH        | 20H         | UOPS_ISSUED.SLOW_LEA                  | Number of slow LEA or similar uops allocated. Such uop has 3 sources (for example, 2 sources + immediate) regardless of whether it is a result of LEA instruction or not. |   |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description  | Comment   |
|------------|-------------|--------------------------------------|--|---|
| 0EH        | 40H         | UOPS_ISSUED.SINGLE_MUL               | Number of multiply packed/scalar single precision uops allocated.  |   |
| 24H        | 21H         | L2_RQSTS.DEMAND_DATA_RD_MISS         | Demand data read requests that missed L2, no rejects.  |   |
| 24H        | 41H         | L2_RQSTS.DEMAND_DATA_RD_HIT          | Demand data read requests that hit L2 cache.   |   |
| 24H        | E1H         | L2_RQSTS.ALL_DEMAND_DATA_RD          | Counts any demand and L1 HW prefetch data load requests to L2.   |   |
| 24H        | 42H         | L2_RQSTS.RFO_HIT                     | Counts the number of store RFO requests that hit the L2 cache.   |   |
| 24H        | 22H         | L2_RQSTS.RFO_MISS                    | Counts the number of store RFO requests that miss the L2 cache.  |   |
| 24H        | E2H         | L2_RQSTS.ALL_RFO                     | Counts all L2 store RFO requests.  |   |
| 24H        | 44H         | L2_RQSTS.CODE_RD_HIT                 | Number of instruction fetches that hit the L2 cache.   |   |
| 24H        | 24H         | L2_RQSTS.CODE_RD_MISS                | Number of instruction fetches that missed the L2 cache.  |   |
| 24H        | 27H         | L2_RQSTS.ALL_DEMAND_MISS             | Demand requests that miss L2 cache.  |   |
| 24H        | E7H         | L2_RQSTS.ALL_DEMAND_REFERENCES       | Demand requests to L2 cache.   |   |
| 24H        | E4H         | L2_RQSTS.ALL_CODE_RD                 | Counts all L2 code requests.   |   |
| 24H        | 50H         | L2_RQSTS.L2_PF_HIT                   | Counts all L2 HW prefetcher requests that hit L2.  |   |
| 24H        | 30H         | L2_RQSTS.L2_PF_MISS                  | Counts all L2 HW prefetcher requests that missed L2.   |   |
| 24H        | F8H         | L2_RQSTS.ALL_PF                      | Counts all L2 HW prefetcher requests.  |   |
| 24H        | 3FH         | L2_RQSTS.MISS                        | All requests that missed L2.   |   |
| 24H        | FFH         | L2_RQSTS.REFERENCES                  | All requests to L2 cache.  |   |
| 27H        | 50H         | L2_DEMAND_RQSTS.WB_HIT               | Not rejected writebacks that hit L2 cache.   |   |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE          | This event counts requests originating from the core that reference a cache line in the last level cache.  | See Table 19-1.                                   |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS               | This event counts each cache miss condition for references to the last level cache.  | See Table 19-1.                                   |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P            | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. | See Table 19-1.                                   |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK     | Increments at the frequency of XCLK (100 MHz) when not halted.   | See Table 19-1.                                   |
| 48H        | 01H         | L1D_PEND_MISS.PENDING                | Increments the number of outstanding L1D misses every cycle. Set Cmask = 1 and Edge = 1 to count occurrences.  | Counter 2 only.<br>Set Cmask = 1 to count cycles. |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK | Miss in all TLB levels causes a page walk of any page size (4K/2M/4M/1G).  |   |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED_4K  | Completed page walks due to store misses in one or more TLB levels of 4K page structure.   |   |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description  | Comment                       |
|------------|-------------|---|--|-------------------------------|
| 49H        | 04H         | DTLB_STORE_MISSES.WALK_COMPLETED_2M_4M      | Completed page walks due to store misses in one or more TLB levels of 2M/4M page structure.        |                               |
| 49H        | 0EH         | DTLB_STORE_MISSES.WALK_COMPLETED            | Completed page walks due to store miss in any TLB levels of any page size (4K/2M/4M/1G).           |                               |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_DURATION             | Cycles PMH is busy with this walk.   |                               |
| 49H        | 20H         | DTLB_STORE_MISSES.STLB_HIT_4K               | Store misses that missed DTLB but hit STLB (4K).   |                               |
| 49H        | 40H         | DTLB_STORE_MISSES.STLB_HIT_2M               | Store misses that missed DTLB but hit STLB (2M).   |                               |
| 49H        | 60H         | DTLB_STORE_MISSES.STLB_HIT                  | Store operations that miss the first TLB level but hit the second and do not cause page walks.     |                               |
| 49H        | 80H         | DTLB_STORE_MISSES.PDE_CACHE_MISS            | DTLB store misses with low part of linear-to-physical address translation missed.                  |                               |
| 4CH        | 01H         | LOAD_HIT_PRE.SW_PF                          | Non-Sw-prefetch load dispatches that hit fill buffer allocated for S/W prefetch.                   |                               |
| 4CH        | 02H         | LOAD_HIT_PRE.HW_PF                          | Non-Sw-prefetch load dispatches that hit fill buffer allocated for H/W prefetch.                   |                               |
| 51H        | 01H         | L1D.REPLACEMENT                             | Counts the number of lines brought into the L1 data cache.   |                               |
| 58H        | 04H         | MOVE_ELIMINATION.INT_NOT_ELIMINATED         | Number of integer move elimination candidate uops that were not eliminated.                        |                               |
| 58H        | 08H         | MOVE_ELIMINATION.SIMD_NOT_ELIMINATED        | Number of SIMD move elimination candidate uops that were not eliminated.                           |                               |
| 58H        | 01H         | MOVE_ELIMINATION.INT_ELIMINATED             | Number of integer move elimination candidate uops that were eliminated.                            |                               |
| 58H        | 02H         | MOVE_ELIMINATION.SIMD_ELIMINATED            | Number of SIMD move elimination candidate uops that were eliminated.                               |                               |
| 5CH        | 01H         | CPL_CYCLES.RING0                            | Unhalted core cycles when the thread is in ring 0.   | Use Edge to count transition. |
| 5CH        | 02H         | CPL_CYCLES.RING123                          | Unhalted core cycles when the thread is not in ring 0.   |                               |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES                      | Cycles the RS is empty for the thread.   |                               |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD | Offcore outstanding demand data read transactions in SQ to uncore. Set Cmask=1 to count cycles.    | Use only when HTT is off.     |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_CODE_RD | Offcore outstanding Demand code Read transactions in SQ to uncore. Set Cmask=1 to count cycles.    | Use only when HTT is off.     |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO     | Offcore outstanding RFO store transactions in SQ to uncore. Set Cmask=1 to count cycles.           | Use only when HTT is off.     |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD    | Offcore outstanding cacheable data read transactions in SQ to uncore. Set Cmask=1 to count cycles. | Use only when HTT is off.     |
| 63H        | 01H         | LOCK_CYCLES.SPLIT_LOCK_UC_LOCK_DURATION     | Cycles in which the L1D and L2 are locked, due to a UC lock or split lock.                         |                               |
| 63H        | 02H         | LOCK_CYCLES.CACHE_LOCK_DURATION             | Cycles in which the L1D is locked.   |                               |
| 79H        | 02H         | IDQ.EMPTY                                   | Counts cycles the IDQ is empty.  |                               |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description  | Comment                           |
|------------|-------------|--|--|-----------------------------------|
| 79H        | 04H         | IDQ.MITE_UOPS                          | Increment each cycle # of uops delivered to IDQ from MITE path. Set Cmask = 1 to count cycles.   | Can combine Umask 04H and 20H.    |
| 79H        | 08H         | IDQ.DSB_UOPS                           | Increment each cycle. # of uops delivered to IDQ from DSB path. Set Cmask = 1 to count cycles.   | Can combine Umask 08H and 10H.    |
| 79H        | 10H         | IDQ.MS_DSB_UOPS                        | Increment each cycle # of uops delivered to IDQ when MS_busy by DSB. Set Cmask = 1 to count cycles. Add Edge=1 to count # of delivery. | Can combine Umask 04H, 08H.       |
| 79H        | 20H         | IDQ.MS_MITE_UOPS                       | Increment each cycle # of uops delivered to IDQ when MS_busy by MITE. Set Cmask = 1 to count cycles.                                   | Can combine Umask 04H, 08H.       |
| 79H        | 30H         | IDQ.MS_UOPS                            | Increment each cycle # of uops delivered to IDQ from MS by either DSB or MITE. Set Cmask = 1 to count cycles.                          | Can combine Umask 04H, 08H.       |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_ANY_UOPS            | Counts cycles DSB is delivered at least one uops. Set Cmask = 1.   |                                   |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_4_UOPS              | Counts cycles DSB is delivered four uops. Set Cmask = 4.   |                                   |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_ANY_UOPS           | Counts cycles MITE is delivered at least one uop. Set Cmask = 1.   |                                   |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_4_UOPS             | Counts cycles MITE is delivered four uops. Set Cmask = 4.  |                                   |
| 79H        | 3CH         | IDQ.MITE_ALL_UOPS                      | # of uops delivered to IDQ from any path.  |                                   |
| 80H        | 02H         | ICACHE.MISSES                          | Number of Instruction Cache, Streaming Buffer and Victim Cache Misses. Includes UC accesses.   |                                   |
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK         | Misses in ITLB that causes a page walk of any page size.   |                                   |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETE_D_4K         | Completed page walks due to misses in ITLB 4K page entries.  |                                   |
| 85H        | 04H         | ITLB_MISSES.WALK_COMPLETE_D_2M_4M      | Completed page walks due to misses in ITLB 2M/4M page entries.   |                                   |
| 85H        | 0EH         | ITLB_MISSES.WALK_COMPLETE_D            | Completed page walks in ITLB of any page size.   |                                   |
| 85H        | 10H         | ITLB_MISSES.WALK_DURATION              | Cycle PMH is busy with a walk.   |                                   |
| 85H        | 20H         | ITLB_MISSES.STLB_HIT_4K                | ITLB misses that hit STLB (4K).  |                                   |
| 85H        | 40H         | ITLB_MISSES.STLB_HIT_2M                | ITLB misses that hit STLB (2M).  |                                   |
| 85H        | 60H         | ITLB_MISSES.STLB_HIT                   | ITLB misses that hit STLB. No page walk.   |                                   |
| 87H        | 01H         | ILD_STALL.LCP                          | Stalls caused by changing prefix length of the instruction.  |                                   |
| 87H        | 04H         | ILD_STALL.IQ_FULL                      | Stall cycles due to IQ is full.  |                                   |
| 88H        | 01H         | BR_INST_EXEC.COND                      | Qualify conditional near branch instructions executed, but not necessarily retired.  | Must combine with umask 40H, 80H. |
| 88H        | 02H         | BR_INST_EXEC.DIRECT_JMP                | Qualify all unconditional near branch instructions excluding calls and indirect branches.  | Must combine with umask 80H.      |
| 88H        | 04H         | BR_INST_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify executed indirect near branch instructions that are not calls or returns.  | Must combine with umask 80H.      |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description   | Comment                           |
|------------|-------------|--|---|-----------------------------------|
| 88H        | 08H         | BR_INST_EXEC.RETURN_NEAR               | Qualify indirect near branches that have a return mnemonic.   | Must combine with umask 80H.      |
| 88H        | 10H         | BR_INST_EXEC.DIRECT_NEAR_CALL          | Qualify unconditional near call branch instructions, excluding non-call branch, executed.                                 | Must combine with umask 80H.      |
| 88H        | 20H         | BR_INST_EXEC.INDIRECT_NEAR_CALL        | Qualify indirect near calls, including both register and memory indirect, executed.                                       | Must combine with umask 80H.      |
| 88H        | 40H         | BR_INST_EXEC.NONTAKEN                  | Qualify non-taken near branches executed.   | Applicable to umask 01H only.     |
| 88H        | 80H         | BR_INST_EXEC.TAKEN                     | Qualify taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                                      |                                   |
| 88H        | FFH         | BR_INST_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 89H        | 01H         | BR_MISP_EXEC.COND                      | Qualify conditional near branch instructions mispredicted.  | Must combine with umask 40H, 80H. |
| 89H        | 04H         | BR_MISP_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify mispredicted indirect near branch instructions that are not calls or returns.                                     | Must combine with umask 80H.      |
| 89H        | 08H         | BR_MISP_EXEC.RETURN_NEAR               | Qualify mispredicted indirect near branches that have a return mnemonic.  | Must combine with umask 80H.      |
| 89H        | 10H         | BR_MISP_EXEC.DIRECT_NEAR_CALL          | Qualify mispredicted unconditional near call branch instructions, excluding non-call branch, executed.                    | Must combine with umask 80H.      |
| 89H        | 20H         | BR_MISP_EXEC.INDIRECT_NEAR_CALL        | Qualify mispredicted indirect near calls, including both register and memory indirect, executed.                          | Must combine with umask 80H.      |
| 89H        | 40H         | BR_MISP_EXEC.NONTAKEN                  | Qualify mispredicted non-taken near branches executed.  | Applicable to umask 01H only.     |
| 89H        | 80H         | BR_MISP_EXEC.TAKEN                     | Qualify mispredicted taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                         |                                   |
| 89H        | FFH         | BR_MISP_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CO<br>RE        | Count issue pipeline slots where no uop was delivered from the front end to the back end when there is no back-end stall. | Use Cmask to qualify uop b/w.     |
| A1H        | 01H         | UOPS_EXECUTED_PORT.PORT_0              | Cycles which a uop is dispatched on port 0 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 02H         | UOPS_EXECUTED_PORT.PORT_1              | Cycles which a uop is dispatched on port 1 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 04H         | UOPS_EXECUTED_PORT.PORT_2              | Cycles which a uop is dispatched on port 2 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 08H         | UOPS_EXECUTED_PORT.PORT_3              | Cycles which a uop is dispatched on port 3 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 10H         | UOPS_EXECUTED_PORT.PORT_4              | Cycles which a uop is dispatched on port 4 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 20H         | UOPS_EXECUTED_PORT.PORT_5              | Cycles which a uop is dispatched on port 5 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 40H         | UOPS_EXECUTED_PORT.PORT_6              | Cycles which a uop is dispatched on port 6 in this thread.  | Set AnyThread to count per core.  |
| A1H        | 80H         | UOPS_EXECUTED_PORT.PORT_7              | Cycles which a uop is dispatched on port 7 in this thread   | Set AnyThread to count per core.  |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description  | Comment                   |
|------------|-------------|-----------------------------------|--|---------------------------|
| A2H        | 01H         | RESOURCE_STALLS.ANY               | Cycles allocation is stalled due to resource related reason.                         |                           |
| A2H        | 04H         | RESOURCE_STALLS.RS                | Cycles stalled due to no eligible RS entry available.                                |                           |
| A2H        | 08H         | RESOURCE_STALLS.SB                | Cycles stalled due to no store buffers available (not including draining from sync). |                           |
| A2H        | 10H         | RESOURCE_STALLS.ROB               | Cycles stalled due to re-order buffer full.  |                           |
| A3H        | 01H         | CYCLE_ACTIVITY.CYCLES_L2_PENDING  | Cycles with pending L2 miss loads. Set Cmask=2 to count cycle.                       | Use only when HTT is off. |
| A3H        | 02H         | CYCLE_ACTIVITY.CYCLES_LDM_PENDING | Cycles with pending memory loads. Set Cmask=2 to count cycle.                        |                           |
| A3H        | 05H         | CYCLE_ACTIVITY.STALLS_L2_PENDING  | Number of loads missed L2.   | Use only when HTT is off. |
| A3H        | 08H         | CYCLE_ACTIVITY.CYCLES_L1D_PENDING | Cycles with pending L1 data cache miss loads. Set Cmask=8 to count cycle.            | PMC2 only.                |
| A3H        | 0CH         | CYCLE_ACTIVITY.STALLS_L1D_PENDING | Execution stalls due to L1 data cache miss loads. Set Cmask=0CH.                     | PMC2 only.                |
| A8H        | 01H         | LSD.UOPS                          | Number of uops delivered by the LSD.   |                           |
| AEH        | 01H         | ITLB.ITLB_FLUSH                   | Counts the number of ITLB flushes, includes 4k/2M/4M pages.                          |                           |
| B0H        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD   | Demand data read requests sent to uncore.  | Use only when HTT is off. |
| B0H        | 02H         | OFFCORE_REQUESTS.DEMAND_CODE_RD   | Demand code read requests sent to uncore.  | Use only when HTT is off. |
| B0H        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO       | Demand RFO read requests sent to uncore, including regular RFOs, locks, ItoM.        | Use only when HTT is off. |
| B0H        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD      | Data read requests sent to uncore (demand and prefetch).                             | Use only when HTT is off. |
| B1H        | 02H         | UOPS_EXECUTED.CORE                | Counts total number of uops to be executed per-core each cycle.                      | Do not need to set ANY.   |
| B7H        | 01H         | OFF_CORE_RESPONSE_0               | See Table 18-28 or Table 18-29.  | Requires MSR 01A6H.       |
| BBH        | 01H         | OFF_CORE_RESPONSE_1               | See Table 18-28 or Table 18-29.  | Requires MSR 01A7H.       |
| BCH        | 11H         | PAGE_WALKER_LOADS.DTLB_L1         | Number of DTLB page walker loads that hit in the L1+FB.                              |                           |
| BCH        | 21H         | PAGE_WALKER_LOADS.ITLB_L1         | Number of ITLB page walker loads that hit in the L1+FB.                              |                           |
| BCH        | 12H         | PAGE_WALKER_LOADS.DTLB_L2         | Number of DTLB page walker loads that hit in the L2.                                 |                           |
| BCH        | 22H         | PAGE_WALKER_LOADS.ITLB_L2         | Number of ITLB page walker loads that hit in the L2.                                 |                           |
| BCH        | 14H         | PAGE_WALKER_LOADS.DTLB_L3         | Number of DTLB page walker loads that hit in the L3.                                 |                           |
| BCH        | 24H         | PAGE_WALKER_LOADS.ITLB_L3         | Number of ITLB page walker loads that hit in the L3.                                 |                           |
| BCH        | 18H         | PAGE_WALKER_LOADS.DTLB_MEMORY     | Number of DTLB page walker loads from memory.  |                           |
| BCH        | 28H         | PAGE_WALKER_LOADS.ITLB_MEMORY     | Number of ITLB page walker loads from memory.  |                           |
| BDH        | 01H         | TLB_FLUSH.DTLB_THREAD             | DTLB flush attempts of the thread-specific entries.                                  |                           |
| BDH        | 20H         | TLB_FLUSH.STLB_ANY                | Count number of STLB flush attempts.   |                           |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment  |
|------------|-------------|--------------------------------|--|--|
| C0H        | 00H         | INST_RETIRED.ANY_P             | Number of instructions at retirement.  | See Table 19-1.  |
| C0H        | 01H         | INST_RETIRED.PREC_DIST         | Precise instruction retired event with HW to reduce effect of PEBS shadow in IP distribution.                                | PMC1 only.   |
| C1H        | 08H         | OTHER_ASSISTS.AVX_TO_SSE       | Number of transitions from AVX-256 to legacy SSE when penalty applicable.  |  |
| C1H        | 10H         | OTHER_ASSISTS.SSE_TO_AVX       | Number of transitions from SSE to AVX-256 when penalty applicable.   |  |
| C1H        | 40H         | OTHER_ASSISTS.ANY_WB_ASSIST    | Number of microcode assists invoked by HW upon uop writeback.  |  |
| C2H        | 01H         | UOPS_RETIRED.ALL               | Counts the number of micro-ops retired. Use Cmask=1 and invert to count active cycles or stalled cycles.                     | Supports PEBS and DataLA; use Any=1 for core granular. |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS      | Counts the number of retirement slots used each cycle.   | Supports PEBS.   |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of machine clears due to memory order conflicts.   |  |
| C3H        | 04H         | MACHINE_CLEARS.SMC             | Number of self-modifying-code machine clears detected.   |  |
| C3H        | 20H         | MACHINE_CLEARS.MASKMOV         | Counts the number of executed AVX masked load operations that refer to an illegal address range with the mask bits set to 0. |  |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Branch instructions at retirement.   | See Table 19-1.  |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL    | Counts the number of conditional branch instructions retired.  | Supports PEBS.   |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL      | Direct and indirect near call instructions retired.  | Supports PEBS.   |
| C4H        | 04H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts the number of branch instructions retired.  | Supports PEBS.   |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETURN    | Counts the number of near return instructions retired.   | Supports PEBS.   |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN      | Counts the number of not taken branch instructions retired.  |  |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKEN     | Number of near taken branches retired.   | Supports PEBS.   |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANCHES   | Number of far branches retired.  |  |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted branch instructions at retirement.  | See Table 19-1.  |
| C5H        | 01H         | BR_MISP_RETIRED.CONDITIONAL    | Mispredicted conditional branch instructions retired.  | Supports PEBS.   |
| C5H        | 04H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted macro branch instructions retired.  | Supports PEBS.   |
| C5H        | 20H         | BR_MISP_RETIRED.NEAR_TAKEN     | Number of near branch instructions retired that were taken but mispredicted.   |  |
| CAH        | 02H         | FP_ASSIST.X87_OUTPUT           | Number of X87 FP assists due to output values.   |  |
| CAH        | 04H         | FP_ASSIST.X87_INPUT            | Number of X87 FP assists due to input values.  |  |



**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description   | Comment                        |
|------------|-------------|--|---|--------------------------------|
| CAH        | 08H         | FP_ASSIST.SIMD_OUTPUT                    | Number of SIMD FP assists due to output values.   |                                |
| CAH        | 10H         | FP_ASSIST.SIMD_INPUT                     | Number of SIMD FP assists due to input values.  |                                |
| CAH        | 1EH         | FP_ASSIST.ANY                            | Cycles with any input/output SSE* or FP assists.  |                                |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSERTS              | Count cases of saving new LBR records by hardware.  |                                |
| CDH        | 01H         | MEM_TRANS_RETIRED.LOAD_LATENCY           | Randomly sampled loads whose latency is above a user defined threshold. A small fraction of the overall loads are sampled due to randomization. | Specify threshold in MSR 3F6H. |
| DOH        | 11H         | MEM_UOPS_RETIRED.STLB_MISSES_LOADS       | Retired load uops that miss the STLB.   | Supports PEBS and DataLA.      |
| DOH        | 12H         | MEM_UOPS_RETIRED.STLB_MISSES_STORES      | Retired store uops that miss the STLB.  | Supports PEBS and DataLA.      |
| DOH        | 21H         | MEM_UOPS_RETIRED.LOCK_LOADS              | Retired load uops with locked access.   | Supports PEBS and DataLA.      |
| DOH        | 41H         | MEM_UOPS_RETIRED.SPLIT_LOADS             | Retired load uops that split across a cacheline boundary.   | Supports PEBS and DataLA.      |
| DOH        | 42H         | MEM_UOPS_RETIRED.SPLIT_STORES            | Retired store uops that split across a cacheline boundary.  | Supports PEBS and DataLA.      |
| DOH        | 81H         | MEM_UOPS_RETIRED.ALL_LOADS               | All retired load uops.  | Supports PEBS and DataLA.      |
| DOH        | 82H         | MEM_UOPS_RETIRED.ALL_STORES              | All retired store uops.   | Supports PEBS and DataLA.      |
| D1H        | 01H         | MEM_LOAD_UOPS_RETIRED.L1_HIT             | Retired load uops with L1 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 02H         | MEM_LOAD_UOPS_RETIRED.L2_HIT             | Retired load uops with L2 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 04H         | MEM_LOAD_UOPS_RETIRED.L3_HIT             | Retired load uops with L3 cache hits as data sources.   | Supports PEBS and DataLA.      |
| D1H        | 08H         | MEM_LOAD_UOPS_RETIRED.L1_MISS            | Retired load uops missed L1 cache as data sources.  | Supports PEBS and DataLA.      |
| D1H        | 10H         | MEM_LOAD_UOPS_RETIRED.L2_MISS            | Retired load uops missed L2. Unknown data source excluded.  | Supports PEBS and DataLA.      |
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRED.L3_MISS            | Retired load uops missed L3. Excludes unknown data source .   | Supports PEBS and DataLA.      |
| D1H        | 40H         | MEM_LOAD_UOPS_RETIRED.HIT_LFB            | Retired load uops which data sources were load uops missed L1 but hit FB due to preceding miss to the same cache line with data not ready.      | Supports PEBS and DataLA.      |
| D2H        | 01H         | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_MISS   | Retired load uops which data sources were L3 hit and cross-core snoop missed in on-pkg core cache.  | Supports PEBS and DataLA.      |
| D2H        | 02H         | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HIT    | Retired load uops which data sources were L3 and cross-core snoop hits in on-pkg core cache.  | Supports PEBS and DataLA.      |
| D2H        | 04H         | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HITM   | Retired load uops which data sources were HitM responses from shared L3.  | Supports PEBS and DataLA.      |
| D2H        | 08H         | MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_NONE   | Retired load uops which data sources were hits in L3 without snoops required.   | Supports PEBS and DataLA.      |
| D3H        | 01H         | MEM_LOAD_UOPS_L3_MISS_RETIRED.LOCAL_DRAM | Retired load uops which data sources missed L3 but serviced from local dram.  | Supports PEBS and DataLA.      |

**Table 19-9. Performance Events in the Processor Core of 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description  | Comment                          |
|------------|-------------|---------------------------|--|----------------------------------|
| E6H        | 1FH         | BACLEAR.S.ANY             | Number of front end re-steers due to BPU misprediction.    |                                  |
| F0H        | 01H         | L2_TRANS.DEMAND_DATA_RD   | Demand data read requests that access L2 cache.            |                                  |
| F0H        | 02H         | L2_TRANS.RFO              | RFO requests that access L2 cache.                         |                                  |
| F0H        | 04H         | L2_TRANS.CODE_RD          | L2 cache accesses when fetching instructions.              |                                  |
| F0H        | 08H         | L2_TRANS.ALL_PF           | Any MLC or L3 HW prefetch accessing L2, including rejects. |                                  |
| F0H        | 10H         | L2_TRANS.L1D_WB           | L1D writebacks that access L2 cache.                       |                                  |
| F0H        | 20H         | L2_TRANS.L2_FILL          | L2 fill requests that access L2 cache.                     |                                  |
| F0H        | 40H         | L2_TRANS.L2_WB            | L2 writebacks that access L2 cache.                        |                                  |
| F0H        | 80H         | L2_TRANS.ALL_REQUESTS     | Transactions accessing L2 pipe.                            |                                  |
| F1H        | 01H         | L2_LINES_IN.I             | L2 cache lines in I state filling L2.                      | Counting does not cover rejects. |
| F1H        | 02H         | L2_LINES_IN.S             | L2 cache lines in S state filling L2.                      | Counting does not cover rejects. |
| F1H        | 04H         | L2_LINES_IN.E             | L2 cache lines in E state filling L2.                      | Counting does not cover rejects. |
| F1H        | 07H         | L2_LINES_IN.ALL           | L2 cache lines filling L2.                                 | Counting does not cover rejects. |
| F2H        | 05H         | L2_LINES_OUT.DEMAND_CLEAN | Clean L2 cache lines evicted by demand.                    |                                  |
| F2H        | 06H         | L2_LINES_OUT.DEMAND_DIRTY | Dirty L2 cache lines evicted by demand.                    |                                  |

**Table 19-10. Intel TSX Performance Events in Processors Based on Haswell Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic                                   | Description  | Comment |
|------------|-------------|---|--|---------|
| 54H        | 01H         | TX_MEM.ABORT_CONFLICT                                 | Number of times a transactional abort was signaled due to a data conflict on a transactionally accessed address.                                     |         |
| 54H        | 02H         | TX_MEM.ABORT_CAPACITY_WRITE                           | Number of times a transactional abort was signaled due to a data capacity limitation for transactional writes.                                       |         |
| 54H        | 04H         | TX_MEM.ABORT_HLE_STORE_TO_ELIDED_LOCK                 | Number of times a HLE transactional region aborted due to a non XRELEASE prefixed instruction writing to an elided lock in the elision buffer.       |         |
| 54H        | 08H         | TX_MEM.ABORT_HLE_ELISION_BUFFER_NOT_EMPTY             | Number of times an HLE transactional execution aborted due to NoAllocatedElisionBuffer being non-zero.   |         |
| 54H        | 10H         | TX_MEM.ABORT_HLE_ELISION_BUFFER_MISMATCH              | Number of times an HLE transactional execution aborted due to XRELEASE lock not satisfying the address and value requirements in the elision buffer. |         |
| 54H        | 20H         | TX_MEM.ABORT_HLE_ELISION_BUFFER_UNSUPPORTED_ALIGNMENT | Number of times an HLE transactional execution aborted due to an unsupported read alignment from the elision buffer.                                 |         |
| 54H        | 40H         | TX_MEM.HLE_ELISION_BUFFER_FULL                        | Number of times HLE lock could not be elided due to ElisionBufferAvailable being zero.   |         |

Table 19-10. Intel TSX Performance Events in Processors Based on Haswell Microarchitecture

| Event Num. | Umask Value | Event Mask Mnemonic      | Description  | Comment              |
|------------|-------------|--------------------------|--|----------------------|
| 5DH        | 01H         | TX_EXEC.MISC1            | Counts the number of times a class of instructions that may cause a transactional abort was executed. Since this is the count of execution, it may not always cause a transactional abort. |                      |
| 5DH        | 02H         | TX_EXEC.MISC2            | Counts the number of times a class of instructions (for example, vzeroupper) that may cause a transactional abort was executed inside a transactional region.                              |                      |
| 5DH        | 04H         | TX_EXEC.MISC3            | Counts the number of times an instruction execution caused the transactional nest count supported to be exceeded.  |                      |
| 5DH        | 08H         | TX_EXEC.MISC4            | Counts the number of times an XBEGIN instruction was executed inside an HLE transactional region.  |                      |
| 5DH        | 10H         | TX_EXEC.MISC5            | Counts the number of times an instruction with HLE-XACQUIRE semantic was executed inside an RTM transactional region.  |                      |
| C8H        | 01H         | HLE_RETIRED.START        | Number of times an HLE execution started.  | IF HLE is supported. |
| C8H        | 02H         | HLE_RETIRED.COMMIT       | Number of times an HLE execution successfully committed.   |                      |
| C8H        | 04H         | HLE_RETIRED.ABORTED      | Number of times an HLE execution aborted due to any reasons (multiple categories may count as one). Supports PEBS.   |                      |
| C8H        | 08H         | HLE_RETIRED.ABORTED_MEM  | Number of times an HLE execution aborted due to various memory events (for example, read/write capacity and conflicts).  |                      |
| C8H        | 10H         | HLE_RETIRED.ABORTED_TIME | Number of times an HLE execution aborted due to uncommon conditions.   |                      |
| C8H        | 20H         | HLE_RETIRED.ABORTED_UNFR | Number of times an HLE execution aborted due to HLE-unfriendly instructions.   |                      |
| C8H        | 40H         | HLE_RETIRED.ABORTED_MEM  | Number of times an HLE execution aborted due to incompatible memory type.  |                      |
| C8H        | 80H         | HLE_RETIRED.ABORTED_EVEN | Number of times an HLE execution aborted due to none of the previous 4 categories (for example, interrupts).   |                      |
| C9H        | 01H         | RTM_RETIRED.START        | Number of times an RTM execution started.  | IF RTM is supported. |
| C9H        | 02H         | RTM_RETIRED.COMMIT       | Number of times an RTM execution successfully committed.   |                      |
| C9H        | 04H         | RTM_RETIRED.ABORTED      | Number of times an RTM execution aborted due to any reasons (multiple categories may count as one). Supports PEBS.   |                      |

**Table 19-10. Intel TSX Performance Events in Processors Based on Haswell Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment              |
|------------|-------------|--------------------------------|---|----------------------|
| C9H        | 08H         | RTM_RETIRED.ABORTED_MEM        | Number of times an RTM execution aborted due to various memory events (for example, read/write capacity and conflicts). | IF RTM is supported. |
| C9H        | 10H         | RTM_RETIRED.ABORTED_TIME R     | Number of times an RTM execution aborted due to uncommon conditions.  |                      |
| C9H        | 20H         | RTM_RETIRED.ABORTED_UNFRIENDLY | Number of times an RTM execution aborted due to HLE-unfriendly instructions.  |                      |
| C9H        | 40H         | RTM_RETIRED.ABORTED_MEMTYPE    | Number of times an RTM execution aborted due to incompatible memory type.   |                      |
| C9H        | 80H         | RTM_RETIRED.ABORTED_EVENTS     | Number of times an RTM execution aborted due to none of the previous 4 categories (for example, interrupt).             |                      |

Model-specific performance monitoring events that are located in the uncore sub-system are implementation specific between different platforms using processors based on Haswell microarchitecture and with different DisplayFamily\_DisplayModel signatures. Processors with CPUID signature of DisplayFamily\_DisplayModel 06\_3CH and 06\_45H support performance events listed in Table 19-11.

**Table 19-11. Uncore Performance Events in the 4th Generation Intel® Core™ Processors**

| Event Num. <sup>1</sup> | Umask Value | Event Mask Mnemonic                   | Description   | Comment  |
|-------------------------|-------------|---------------------------------------|---|--|
| 22H                     | 01H         | UNC_CBO_XSNP_RESPONSE.MISS            | A snoop misses in some processor core.  | Must combine with one of the umask values of 20H, 40H, 80H.      |
| 22H                     | 02H         | UNC_CBO_XSNP_RESPONSE.INVALID         | A snoop invalidates a non-modified line in some processor core.   |  |
| 22H                     | 04H         | UNC_CBO_XSNP_RESPONSE.HIT             | A snoop hits a non-modified line in some processor core.  |  |
| 22H                     | 08H         | UNC_CBO_XSNP_RESPONSE.HITM            | A snoop hits a modified line in some processor core.  |  |
| 22H                     | 10H         | UNC_CBO_XSNP_RESPONSE.INVALID_M       | A snoop invalidates a modified line in some processor core.   |  |
| 22H                     | 20H         | UNC_CBO_XSNP_RESPONSE.EXTERNAL_FILTER | Filter on cross-core snoops initiated by this Cbox due to external snoop request.                                 | Must combine with at least one of 01H, 02H, 04H, 08H, 10H.       |
| 22H                     | 40H         | UNC_CBO_XSNP_RESPONSE.CORE_FILTER     | Filter on cross-core snoops initiated by this Cbox due to processor core memory request.                          |  |
| 22H                     | 80H         | UNC_CBO_XSNP_RESPONSE.EVICTION_FILTER | Filter on cross-core snoops initiated by this Cbox due to L3 eviction.  |  |
| 34H                     | 01H         | UNC_CBO_CACHE_LOOKUP.M                | L3 lookup request that access cache and found line in M-state.  | Must combine with one of the umask values of 10H, 20H, 40H, 80H. |
| 34H                     | 06H         | UNC_CBO_CACHE_LOOKUP.E                | L3 lookup request that access cache and found line in E or S state.   |  |
| 34H                     | 08H         | UNC_CBO_CACHE_LOOKUP.I                | L3 lookup request that access cache and found line in I-state.  |  |
| 34H                     | 10H         | UNC_CBO_CACHE_LOOKUP.READ_FILTER      | Filter on processor core initiated cacheable read requests. Must combine with at least one of 01H, 02H, 04H, 08H. |  |

**Table 19-11. Uncore Performance Events in the 4th Generation Intel® Core™ Processors (Contd.)**

| Event Num. <sup>1</sup> | Umask Value | Event Mask Mnemonic                     | Description  | Comment         |
|-------------------------|-------------|---|--|-----------------|
| 34H                     | 20H         | UNC_CBO_CACHE_LOOKUP.WRITE_FILTER       | Filter on processor core initiated cacheable write requests. Must combine with at least one of 01H, 02H, 04H, 08H.   |                 |
| 34H                     | 40H         | UNC_CBO_CACHE_LOOKUP.EXTSNP_FILTER      | Filter on external snoop requests. Must combine with at least one of 01H, 02H, 04H, 08H.   |                 |
| 34H                     | 80H         | UNC_CBO_CACHE_LOOKUP.ANY_REQUEST_FILTER | Filter on any IRQ or IPQ initiated requests including uncacheable, non-coherent requests. Must combine with at least one of 01H, 02H, 04H, 08H.  |                 |
| 80H                     | 01H         | UNC_ARB_TRK_OCCUPANCY.ALL               | Counts cycles weighted by the number of requests waiting for data returning from the memory controller. Accounts for coherent and non-coherent requests initiated by IA cores, processor graphic units, or L3. | Counter 0 only. |
| 81H                     | 01H         | UNC_ARB_TRK_REQUEST.ALL                 | Counts the number of coherent and in-coherent requests initiated by IA cores, processor graphic units, or L3.  |                 |
| 81H                     | 20H         | UNC_ARB_TRK_REQUEST.WRITES              | Counts the number of allocated write entries, include full, partial, and L3 evictions.   |                 |
| 81H                     | 80H         | UNC_ARB_TRK_REQUEST.EVICTIONS           | Counts the number of L3 evictions allocated.   |                 |
| 83H                     | 01H         | UNC_ARB_COH_TRK_OCCUPANCY.ALL           | Cycles weighted by number of requests pending in Coherency Tracker.  | Counter 0 only. |
| 84H                     | 01H         | UNC_ARB_COH_TRK_REQUEST.ALL             | Number of requests allocated in Coherency Tracker.   |                 |

**NOTES:**

1. The uncore events must be programmed using MSRs located in specific performance monitoring units in the uncore. UNC\_CBO\* events are supported using MSR\_UNC\_CBO\* MSRs; UNC\_ARB\* events are supported using MSR\_UNC\_ARB\*MSRs.

**19.6.1 Performance Monitoring Events in the Processor Core of Intel Xeon Processor E5 v3 Family**

Model-specific performance monitoring events in the processor core that are applicable only to Intel Xeon processor E5 v3 family based on the Haswell-E microarchitecture, with CPUID signature of DisplayFamily\_DisplayModel 06\_3FH, are listed in Table 19-12. The performance events listed in Table 19-9 and Table 19-10 also apply Intel Xeon processor E5 v3 family, except that the OFF\_CORE\_RESPONSE\_x event listed in Table 19-9 should reference Table 18-30.

Uncore performance monitoring events for Intel Xeon Processor E5 v3 families are described in “Intel® Xeon® Processor E5 v3 Uncore Performance Monitoring Programming Reference Manual”.

**Table 19-12. Performance Events Applicable only to the Processor Core of Intel® Xeon® Processor E5 v3 Family**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description   | Comment        |
|------------|-------------|---|---|----------------|
| D3H        | 04H         | MEM_LOAD_UOPS_L3_MISS_RETIRED.REMOTE_DRAM | Retired load uops whose data sources were remote DRAM (snoop not needed, Snoop Miss). | Supports PEBS. |
| D3H        | 10H         | MEM_LOAD_UOPS_L3_MISS_RETIRED.REMOTE_HITM | Retired load uops whose data sources were remote cache HITM.                          | Supports PEBS. |
| D3H        | 20H         | MEM_LOAD_UOPS_L3_MISS_RETIRED.REMOTE_FWD  | Retired load uops whose data sources were forwards from a remote cache.               | Supports PEBS. |

## 19.7 PERFORMANCE MONITORING EVENTS FOR 3RD GENERATION INTEL® CORE™ PROCESSORS

3rd generation Intel® Core™ processors and Intel Xeon processor E3-1200 v2 product family are based on Intel microarchitecture code name Ivy Bridge. They support architectural performance monitoring events listed in Table 19-1. Model-specific performance monitoring events in the processor core are listed in Table 19-13. The events in Table 19-13 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_3AH. Fixed counters in the core PMU support the architecture events defined in Table 19-24.

Additional information on event specifics (e.g. derivative events using specific IA32\_PERFVTSELx modifiers, limitations, special notes and recommendations) can be found at <https://software.intel.com/en-us/forums/software-tuning-performance-optimization-platform-monitoring>.

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description  | Comment   |
|------------|-------------|---|--|---|
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD                   | Loads blocked by overlapping with store buffer that cannot be forwarded.   |   |
| 03H        | 08H         | LD_BLOCKS.NO_SR                           | The number of times that split load operations are temporarily blocked because all resources for handling the split accesses are in use.               |   |
| 05H        | 01H         | MISALIGN_MEM_REF.LOADS                    | Speculative cache-line split load uops dispatched to L1D.  |   |
| 05H        | 02H         | MISALIGN_MEM_REF.STORES                   | Speculative cache-line split Store-address uops dispatched to L1D.   |   |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS           | False dependencies in MOB due to partial compare on address.   |   |
| 08H        | 81H         | DTLB_LOAD_MISSES.MISS_CAUSE_S_A_WALK      | Misses in all TLB levels that cause a page walk of any page size from demand loads.  |   |
| 08H        | 82H         | DTLB_LOAD_MISSES.WALK_COMPLETED           | Misses in all TLB levels that caused page walk completed of any size by demand loads.  |   |
| 08H        | 84H         | DTLB_LOAD_MISSES.WALK_DURATION            | Cycle PMH is busy with a walk due to demand loads.   |   |
| 08H        | 88H         | DTLB_LOAD_MISSES.LARGE_PAGE_WALK_DURATION | Page walk for a large page completed for Demand load.  |   |
| 0EH        | 01H         | UOPS_ISSUED.ANY                           | Increments each cycle the # of Uops issued by the RAT to RS. Set Cmask = 1, Inv = 1, Any= 1 to count stalled cycles of this core.                      | Set Cmask = 1, Inv = 1 to count stalled cycles. |
| 0EH        | 10H         | UOPS_ISSUED.FLAGS_MERGE                   | Number of flags-merge uops allocated. Such uops adds delay.  |   |
| 0EH        | 20H         | UOPS_ISSUED.SLOW_LEA                      | Number of slow LEA or similar uops allocated. Such uop has 3 sources (e.g. 2 sources + immediate) regardless if as a result of LEA instruction or not. |   |
| 0EH        | 40H         | UOPS_ISSUED.SINGLE_MUL                    | Number of multiply packed/scalar single precision uops allocated.  |   |
| 10H        | 01H         | FP_COMP_OPS_EXE.X87                       | Counts number of X87 uops executed.  |   |
| 10H        | 10H         | FP_COMP_OPS_EXE.SSE_FP_PACKED_DOUBLE      | Counts number of SSE* or AVX-128 double precision FP packed uops executed.   |   |
| 10H        | 20H         | FP_COMP_OPS_EXE.SSE_FP_SCALAR_SINGLE      | Counts number of SSE* or AVX-128 single precision FP scalar uops executed.   |   |
| 10H        | 40H         | FP_COMP_OPS_EXE.SSE_PACKED_SINGLE         | Counts number of SSE* or AVX-128 single precision FP packed uops executed.   |   |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description  | Comment         |
|------------|-------------|-----------------------------------|--|-----------------|
| 10H        | 80H         | FP_COMP_OPS_EXE.SSE_SCALAR_DOUBLE | Counts number of SSE* or AVX-128 double precision FP scalar uops executed.   |                 |
| 11H        | 01H         | SIMD_FP_256.PACKED_SINGLE         | Counts 256-bit packed single-precision floating-point instructions.  |                 |
| 11H        | 02H         | SIMD_FP_256.PACKED_DOUBLE         | Counts 256-bit packed double-precision floating-point instructions.  |                 |
| 14H        | 01H         | ARITH.FPU_DIV_ACTIVE              | Cycles that the divider is active, includes INT and FP. Set 'edge =1, cmask=1' to count the number of divides.   |                 |
| 24H        | 01H         | L2_RQSTS.DEMAND_DATA_RD_HIT       | Demand Data Read requests that hit L2 cache.   |                 |
| 24H        | 03H         | L2_RQSTS.ALL_DEMAND_DATA_RD       | Counts any demand and L1 HW prefetch data load requests to L2.   |                 |
| 24H        | 04H         | L2_RQSTS.RFO_HITS                 | Counts the number of store RFO requests that hit the L2 cache.   |                 |
| 24H        | 08H         | L2_RQSTS.RFO_MISS                 | Counts the number of store RFO requests that miss the L2 cache.  |                 |
| 24H        | 0CH         | L2_RQSTS.ALL_RFO                  | Counts all L2 store RFO requests.  |                 |
| 24H        | 10H         | L2_RQSTS.CODE_RD_HIT              | Number of instruction fetches that hit the L2 cache.   |                 |
| 24H        | 20H         | L2_RQSTS.CODE_RD_MISS             | Number of instruction fetches that missed the L2 cache.  |                 |
| 24H        | 30H         | L2_RQSTS.ALL_CODE_RD              | Counts all L2 code requests.   |                 |
| 24H        | 40H         | L2_RQSTS.PF_HIT                   | Counts all L2 HW prefetcher requests that hit L2.  |                 |
| 24H        | 80H         | L2_RQSTS.PF_MISS                  | Counts all L2 HW prefetcher requests that missed L2.   |                 |
| 24H        | C0H         | L2_RQSTS.ALL_PF                   | Counts all L2 HW prefetcher requests.  |                 |
| 27H        | 01H         | L2_STORE_LOCK_RQSTS.MISS          | RFOs that miss cache lines.  |                 |
| 27H        | 08H         | L2_STORE_LOCK_RQSTS.HIT_M         | RFOs that hit cache lines in M state.  |                 |
| 27H        | 0FH         | L2_STORE_LOCK_RQSTS.ALL           | RFOs that access cache lines in any state.   |                 |
| 28H        | 01H         | L2_L1D_WB_RQSTS.MISS              | Not rejected writebacks that missed LLC.   |                 |
| 28H        | 04H         | L2_L1D_WB_RQSTS.HIT_E             | Not rejected writebacks from L1D to L2 cache lines in E state.   |                 |
| 28H        | 08H         | L2_L1D_WB_RQSTS.HIT_M             | Not rejected writebacks from L1D to L2 cache lines in M state.   |                 |
| 28H        | 0FH         | L2_L1D_WB_RQSTS.ALL               | Not rejected writebacks from L1D to L2 cache lines in any state.   |                 |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE       | This event counts requests originating from the core that reference a cache line in the last level cache.  | See Table 19-1  |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS            | This event counts each cache miss condition for references to the last level cache.  | See Table 19-1  |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P         | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. | See Table 19-1. |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description  | Comment                                      |
|------------|-------------|---|--|--|
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK            | Increments at the frequency of XCLK (100 MHz) when not halted.   | See Table 19-1.                              |
| 48H        | 01H         | L1D_PEND_MISS.PENDING                       | Increments the number of outstanding L1D misses every cycle. Set Cmask = 1 and Edge =1 to count occurrences. | PMC2 only;<br>Set Cmask = 1 to count cycles. |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK        | Miss in all TLB levels causes a page walk of any page size (4K/2M/4M/1G).                                    |  |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED            | Miss in all TLB levels causes a page walk that completes of any page size (4K/2M/4M/1G).                     |  |
| 49H        | 04H         | DTLB_STORE_MISSES.WALK_DURATION             | Cycles PMH is busy with this walk.   |  |
| 49H        | 10H         | DTLB_STORE_MISSES.STLB_HIT                  | Store operations that miss the first TLB level but hit the second and do not cause page walks.               |  |
| 4CH        | 01H         | LOAD_HIT_PRE.SW_PF                          | Non-Sw-prefetch load dispatches that hit fill buffer allocated for S/W prefetch.                             |  |
| 4CH        | 02H         | LOAD_HIT_PRE.HW_PF                          | Non-Sw-prefetch load dispatches that hit fill buffer allocated for H/W prefetch.                             |  |
| 51H        | 01H         | L1D.REPLACEMENT                             | Counts the number of lines brought into the L1 data cache.   |  |
| 58H        | 04H         | MOVE_ELIMINATION.INT_NOT_ELIMINATED         | Number of integer Move Elimination candidate uops that were not eliminated.                                  |  |
| 58H        | 08H         | MOVE_ELIMINATION.SIMD_NOT_ELIMINATED        | Number of SIMD Move Elimination candidate uops that were not eliminated.                                     |  |
| 58H        | 01H         | MOVE_ELIMINATION.INT_ELIMINATED             | Number of integer Move Elimination candidate uops that were eliminated.                                      |  |
| 58H        | 02H         | MOVE_ELIMINATION.SIMD_ELIMINATED            | Number of SIMD Move Elimination candidate uops that were eliminated.   |  |
| 5CH        | 01H         | CPL_CYCLES.RING0                            | Unhalted core cycles when the thread is in ring 0.   | Use Edge to count transition.                |
| 5CH        | 02H         | CPL_CYCLES.RING123                          | Unhalted core cycles when the thread is not in ring 0.   |  |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES                      | Cycles the RS is empty for the thread.   |  |
| 5FH        | 04H         | DTLB_LOAD_MISSES.STLB_HIT                   | Counts load operations that missed 1st level DTLB but hit the 2nd level.                                     |  |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD | Offcore outstanding Demand Data Read transactions in SQ to uncore. Set Cmask=1 to count cycles.              |  |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_CODE_RD | Offcore outstanding Demand Code Read transactions in SQ to uncore. Set Cmask=1 to count cycles.              |  |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO     | Offcore outstanding RFO store transactions in SQ to uncore. Set Cmask=1 to count cycles.                     |  |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD    | Offcore outstanding cacheable data read transactions in SQ to uncore. Set Cmask=1 to count cycles.           |  |
| 63H        | 01H         | LOCK_CYCLES.SPLIT_LOCK_UC_LOCK_DURATION     | Cycles in which the L1D and L2 are locked, due to a UC lock or split lock.                                   |  |



**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description  | Comment                           |
|------------|-------------|--|--|-----------------------------------|
| 63H        | 02H         | LOCK_CYCLES.CACHE_LOCK_DURATION        | Cycles in which the L1D is locked.   |                                   |
| 79H        | 02H         | IDQ.EMPTY                              | Counts cycles the IDQ is empty.  |                                   |
| 79H        | 04H         | IDQ.MITE_UOPS                          | Increment each cycle # of uops delivered to IDQ from MITE path. Set Cmask = 1 to count cycles.   | Can combine Umask 04H and 20H.    |
| 79H        | 08H         | IDQ.DSB_UOPS                           | Increment each cycle. # of uops delivered to IDQ from DSB path. Set Cmask = 1 to count cycles.   | Can combine Umask 08H and 10H.    |
| 79H        | 10H         | IDQ.MS_DSB_UOPS                        | Increment each cycle # of uops delivered to IDQ when MS_busy by DSB. Set Cmask = 1 to count cycles. Add Edge=1 to count # of delivery. | Can combine Umask 04H, 08H.       |
| 79H        | 20H         | IDQ.MS_MITE_UOPS                       | Increment each cycle # of uops delivered to IDQ when MS_busy by MITE. Set Cmask = 1 to count cycles.                                   | Can combine Umask 04H, 08H.       |
| 79H        | 30H         | IDQ.MS_UOPS                            | Increment each cycle # of uops delivered to IDQ from MS by either DSB or MITE. Set Cmask = 1 to count cycles.                          | Can combine Umask 04H, 08H.       |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_ANY_UOPS            | Counts cycles DSB is delivered at least one uops. Set Cmask = 1.   |                                   |
| 79H        | 18H         | IDQ.ALL_DSB_CYCLES_4_UOPS              | Counts cycles DSB is delivered four uops. Set Cmask = 4.   |                                   |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_ANY_UOPS           | Counts cycles MITE is delivered at least one uops. Set Cmask = 1.  |                                   |
| 79H        | 24H         | IDQ.ALL_MITE_CYCLES_4_UOPS             | Counts cycles MITE is delivered four uops. Set Cmask = 4.  |                                   |
| 79H        | 3CH         | IDQ.MITE_ALL_UOPS                      | # of uops delivered to IDQ from any path.  |                                   |
| 80H        | 04H         | ICACHE.IFETCH_STALL                    | Cycles where a code-fetch stalled due to L1 instruction-cache miss or an iTLB miss.  |                                   |
| 80H        | 02H         | ICACHE.MISSES                          | Number of Instruction Cache, Streaming Buffer and Victim Cache Misses. Includes UC accesses.   |                                   |
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK         | Misses in all ITLB levels that cause page walks.   |                                   |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETED             | Misses in all ITLB levels that cause completed page walks.   |                                   |
| 85H        | 04H         | ITLB_MISSES.WALK_DURATION              | Cycle PMH is busy with a walk.   |                                   |
| 85H        | 10H         | ITLB_MISSES.STLB_HIT                   | Number of cache load STLB hits. No page walk.  |                                   |
| 87H        | 01H         | ILD_STALL.LCP                          | Stalls caused by changing prefix length of the instruction.  |                                   |
| 87H        | 04H         | ILD_STALL.IQ_FULL                      | Stall cycles due to IQ is full.  |                                   |
| 88H        | 01H         | BR_INST_EXEC.COND                      | Qualify conditional near branch instructions executed, but not necessarily retired.  | Must combine with umask 40H, 80H. |
| 88H        | 02H         | BR_INST_EXEC.DIRECT_JMP                | Qualify all unconditional near branch instructions excluding calls and indirect branches.  | Must combine with umask 80H.      |
| 88H        | 04H         | BR_INST_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify executed indirect near branch instructions that are not calls or returns.  | Must combine with umask 80H.      |
| 88H        | 08H         | BR_INST_EXEC.RETURN_NEAR               | Qualify indirect near branches that have a return mnemonic.  | Must combine with umask 80H.      |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description   | Comment                           |
|------------|-------------|--|---|-----------------------------------|
| 88H        | 10H         | BR_INST_EXEC.DIRECT_NEAR_CALL          | Qualify unconditional near call branch instructions, excluding non-call branch, executed.                                 | Must combine with umask 80H.      |
| 88H        | 20H         | BR_INST_EXEC.INDIRECT_NEAR_CALL        | Qualify indirect near calls, including both register and memory indirect, executed.                                       | Must combine with umask 80H.      |
| 88H        | 40H         | BR_INST_EXEC.NONTAKEN                  | Qualify non-taken near branches executed.   | Applicable to umask 01H only.     |
| 88H        | 80H         | BR_INST_EXEC.TAKEN                     | Qualify taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                                      |                                   |
| 88H        | FFH         | BR_INST_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 89H        | 01H         | BR_MISP_EXEC.COND                      | Qualify conditional near branch instructions mispredicted.  | Must combine with umask 40H, 80H. |
| 89H        | 04H         | BR_MISP_EXEC.INDIRECT_JMP_NON_CALL_RET | Qualify mispredicted indirect near branch instructions that are not calls or returns.                                     | Must combine with umask 80H.      |
| 89H        | 08H         | BR_MISP_EXEC.RETURN_NEAR               | Qualify mispredicted indirect near branches that have a return mnemonic.  | Must combine with umask 80H.      |
| 89H        | 10H         | BR_MISP_EXEC.DIRECT_NEAR_CALL          | Qualify mispredicted unconditional near call branch instructions, excluding non-call branch, executed.                    | Must combine with umask 80H.      |
| 89H        | 20H         | BR_MISP_EXEC.INDIRECT_NEAR_CALL        | Qualify mispredicted indirect near calls, including both register and memory indirect, executed.                          | Must combine with umask 80H.      |
| 89H        | 40H         | BR_MISP_EXEC.NONTAKEN                  | Qualify mispredicted non-taken near branches executed.  | Applicable to umask 01H only.     |
| 89H        | 80H         | BR_MISP_EXEC.TAKEN                     | Qualify mispredicted taken near branches executed. Must combine with 01H,02H, 04H, 08H, 10H, 20H.                         |                                   |
| 89H        | FFH         | BR_MISP_EXEC.ALL_BRANCHES              | Counts all near executed branches (not necessarily retired).  |                                   |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CORE            | Count issue pipeline slots where no uop was delivered from the front end to the back end when there is no back-end stall. | Use Cmask to qualify uop b/w.     |
| A1H        | 01H         | UOPS_DISPATCHED_PORT.PORT_0            | Cycles which a Uop is dispatched on port 0.   |                                   |
| A1H        | 02H         | UOPS_DISPATCHED_PORT.PORT_1            | Cycles which a Uop is dispatched on port 1.   |                                   |
| A1H        | 0CH         | UOPS_DISPATCHED_PORT.PORT_2            | Cycles which a Uop is dispatched on port 2.   |                                   |
| A1H        | 30H         | UOPS_DISPATCHED_PORT.PORT_3            | Cycles which a Uop is dispatched on port 3.   |                                   |
| A1H        | 40H         | UOPS_DISPATCHED_PORT.PORT_4            | Cycles which a Uop is dispatched on port 4.   |                                   |
| A1H        | 80H         | UOPS_DISPATCHED_PORT.PORT_5            | Cycles which a Uop is dispatched on port 5.   |                                   |
| A2H        | 01H         | RESOURCE_STALLS.ANY                    | Cycles Allocation is stalled due to Resource Related reason.  |                                   |
| A2H        | 04H         | RESOURCE_STALLS.RS                     | Cycles stalled due to no eligible RS entry available.   |                                   |
| A2H        | 08H         | RESOURCE_STALLS.SB                     | Cycles stalled due to no store buffers available (not including draining form sync).                                      |                                   |
| A2H        | 10H         | RESOURCE_STALLS.ROB                    | Cycles stalled due to re-order buffer full.   |                                   |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description   | Comment  |
|------------|-------------|-----------------------------------|---|--|
| A3H        | 01H         | CYCLE_ACTIVITY.CYCLES_L2_PENDING  | Cycles with pending L2 miss loads. Set AnyThread to count per core.   |  |
| A3H        | 02H         | CYCLE_ACTIVITY.CYCLES_LDM_PENDING | Cycles with pending memory loads. Set AnyThread to count per core.  | Restricted to counters 0-3 when HTT is disabled. |
| A3H        | 04H         | CYCLE_ACTIVITY.CYCLES_NO_EXECUTE  | Cycles of dispatch stalls. Set AnyThread to count per core.   | Restricted to counters 0-3 when HTT is disabled. |
| A3H        | 05H         | CYCLE_ACTIVITY.STALLS_L2_PENDING  | Number of loads missed L2.  | Restricted to counters 0-3 when HTT is disabled. |
| A3H        | 06H         | CYCLE_ACTIVITY.STALLS_LDM_PENDING |   | Restricted to counters 0-3 when HTT is disabled. |
| A3H        | 08H         | CYCLE_ACTIVITY.CYCLES_L1D_PENDING | Cycles with pending L1 cache miss loads. Set AnyThread to count per core.                                       | PMC2 only.                                       |
| A3H        | 0CH         | CYCLE_ACTIVITY.STALLS_L1D_PENDING | Execution stalls due to L1 data cache miss loads. Set Cmask=0CH.  | PMC2 only.                                       |
| A8H        | 01H         | LSD.UOPS                          | Number of Uops delivered by the LSD.  |  |
| ABH        | 01H         | DSB2MITE_SWITCHES.COUNT           | Number of DSB to MITE switches.   |  |
| ABH        | 02H         | DSB2MITE_SWITCHES.PENALTY_CYCLES  | Cycles DSB to MITE switches caused delay.   |  |
| ACH        | 08H         | DSB_FILL.EXCEED_DSB_LINES         | DSB Fill encountered > 3 DSB lines.   |  |
| AEH        | 01H         | ITLB.ITLB_FLUSH                   | Counts the number of ITLB flushes, includes 4k/2M/4M pages.   |  |
| B0H        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD   | Demand data read requests sent to uncore.   |  |
| B0H        | 02H         | OFFCORE_REQUESTS.DEMAND_CODE_RD   | Demand code read requests sent to uncore.   |  |
| B0H        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO       | Demand RFO read requests sent to uncore, including regular RFOs, locks, ItoM.                                   |  |
| B0H        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD      | Data read requests sent to uncore (demand and prefetch).  |  |
| B1H        | 01H         | UOPS_EXECUTED.THREAD              | Counts total number of uops to be executed per-thread each cycle. Set Cmask = 1, INV = 1 to count stall cycles. |  |
| B1H        | 02H         | UOPS_EXECUTED.CORE                | Counts total number of uops to be executed per-core each cycle.   | Do not need to set ANY.                          |
| B7H        | 01H         | OFFCORE_RESPONSE_0                | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A6H.                              |
| BBH        | 01H         | OFFCORE_RESPONSE_1                | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A7H.                              |
| BDH        | 01H         | TLB_FLUSH.DTLB_THREAD             | DTLB flush attempts of the thread-specific entries.   |  |
| BDH        | 20H         | TLB_FLUSH.STLB_ANY                | Count number of STLB flush attempts.  |  |
| COH        | 00H         | INST_RETIRED.ANY_P                | Number of instructions at retirement.   | See Table 19-1.                                  |
| COH        | 01H         | INST_RETIRED.PREC_DIST            | Precise instruction retired event with HW to reduce effect of PEBS shadow in IP distribution.                   | PMC1 only.                                       |
| C1H        | 08H         | OTHER_ASSISTS.AVX_STORE           | Number of assists associated with 256-bit AVX store operations.   |  |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment                                     |
|------------|-------------|--------------------------------|--|---|
| C1H        | 10H         | OTHER_ASSISTS.AVX_TO_SSE       | Number of transitions from AVX-256 to legacy SSE when penalty applicable.  |   |
| C1H        | 20H         | OTHER_ASSISTS.SSE_TO_AVX       | Number of transitions from SSE to AVX-256 when penalty applicable.   |   |
| C1H        | 80H         | OTHER_ASSISTS.WB               | Number of times microcode assist is invoked by hardware upon uop writeback.  |   |
| C2H        | 01H         | UOPS_RETIRED.ALL               | Counts the number of micro-ops retired, Use cmask=1 and invert to count active cycles or stalled cycles.                     | Supports PEBS, use Any=1 for core granular. |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS      | Counts the number of retirement slots used each cycle.   | Supports PEBS.                              |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of machine clears due to memory order conflicts.   |   |
| C3H        | 04H         | MACHINE_CLEARS.SMC             | Number of self-modifying-code machine clears detected.   |   |
| C3H        | 20H         | MACHINE_CLEARS.MASKMOV         | Counts the number of executed AVX masked load operations that refer to an illegal address range with the mask bits set to 0. |   |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Branch instructions at retirement.   | See Table 19-1.                             |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL    | Counts the number of conditional branch instructions retired.  | Supports PEBS.                              |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL      | Direct and indirect near call instructions retired.  | Supports PEBS.                              |
| C4H        | 04H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts the number of branch instructions retired.  | Supports PEBS.                              |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETURN    | Counts the number of near return instructions retired.   | Supports PEBS.                              |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN      | Counts the number of not taken branch instructions retired.  | Supports PEBS.                              |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKEN     | Number of near taken branches retired.   | Supports PEBS.                              |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANCH     | Number of far branches retired.  | Supports PEBS.                              |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted branch instructions at retirement.  | See Table 19-1.                             |
| C5H        | 01H         | BR_MISP_RETIRED.CONDITIONAL    | Mispredicted conditional branch instructions retired.  | Supports PEBS.                              |
| C5H        | 04H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted macro branch instructions retired.  | Supports PEBS.                              |
| C5H        | 20H         | BR_MISP_RETIRED.NEAR_TAKEN     | Mispredicted taken branch instructions retired.  | Supports PEBS.                              |
| CAH        | 02H         | FP_ASSIST.X87_OUTPUT           | Number of X87 FP assists due to output values.   | Supports PEBS.                              |
| CAH        | 04H         | FP_ASSIST.X87_INPUT            | Number of X87 FP assists due to input values.  | Supports PEBS.                              |
| CAH        | 08H         | FP_ASSIST.SIMD_OUTPUT          | Number of SIMD FP assists due to output values.  | Supports PEBS.                              |
| CAH        | 10H         | FP_ASSIST.SIMD_INPUT           | Number of SIMD FP assists due to input values.   |   |
| CAH        | 1EH         | FP_ASSIST.ANY                  | Cycles with any input/output SSE* or FP assists.   |   |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSERTS    | Count cases of saving new LBR records by hardware.   |   |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description   | Comment   |
|------------|-------------|---|---|---|
| CDH        | 01H         | MEM_TRANS_RETIRED.LOAD_LATENCY            | Randomly sampled loads whose latency is above a user defined threshold. A small fraction of the overall loads are sampled due to randomization. | Specify threshold in MSR 3F6H. PMC 3 only.                      |
| CDH        | 02H         | MEM_TRANS_RETIRED.PRECISE_STORE           | Sample stores and collect precise store operation via PEBS record. PMC3 only.   | See Section 18.3.4.4.3.   |
| DOH        | 11H         | MEM_UOPS_RETIRED.STLB_MISS_LOADS          | Retired load uops that miss the STLB.   | Supports PEBS.  |
| DOH        | 12H         | MEM_UOPS_RETIRED.STLB_MISS_STORES         | Retired store uops that miss the STLB.  | Supports PEBS.  |
| DOH        | 21H         | MEM_UOPS_RETIRED.LOCK_LOADS               | Retired load uops with locked access.   | Supports PEBS.  |
| DOH        | 41H         | MEM_UOPS_RETIRED.SPLIT_LOADS              | Retired load uops that split across a cacheline boundary.   | Supports PEBS.  |
| DOH        | 42H         | MEM_UOPS_RETIRED.SPLIT_STORES             | Retired store uops that split across a cacheline boundary.  | Supports PEBS.  |
| DOH        | 81H         | MEM_UOPS_RETIRED.ALL_LOADS                | All retired load uops.  | Supports PEBS.  |
| DOH        | 82H         | MEM_UOPS_RETIRED.ALL_STORES               | All retired store uops.   | Supports PEBS.  |
| D1H        | 01H         | MEM_LOAD_UOPS_RETIRED.L1_HIT              | Retired load uops with L1 cache hits as data sources.   | Supports PEBS.  |
| D1H        | 02H         | MEM_LOAD_UOPS_RETIRED.L2_HIT              | Retired load uops with L2 cache hits as data sources.   | Supports PEBS.  |
| D1H        | 04H         | MEM_LOAD_UOPS_RETIRED.LLC_HIT             | Retired load uops whose data source was LLC hit with no snoop required.   | Supports PEBS.  |
| D1H        | 08H         | MEM_LOAD_UOPS_RETIRED.L1_MISS             | Retired load uops whose data source followed an L1 miss.  | Supports PEBS.  |
| D1H        | 10H         | MEM_LOAD_UOPS_RETIRED.L2_MISS             | Retired load uops that missed L2, excluding unknown sources.  | Supports PEBS.  |
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRED.LLC_MISS            | Retired load uops whose data source is LLC miss.  | Supports PEBS. Restricted to counters 0-3 when HTT is disabled. |
| D1H        | 40H         | MEM_LOAD_UOPS_RETIRED.HIT_LFB             | Retired load uops which data sources were load uops missed L1 but hit FB due to preceding miss to the same cache line with data not ready.      | Supports PEBS.  |
| D2H        | 01H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_MISS   | Retired load uops whose data source was an on-package core cache LLC hit and cross-core snoop missed.   | Supports PEBS.  |
| D2H        | 02H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HIT    | Retired load uops whose data source was an on-package LLC hit and cross-core snoop hits.  | Supports PEBS.  |
| D2H        | 04H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM   | Retired load uops whose data source was an on-package core cache with HitM responses.   | Supports PEBS.  |
| D2H        | 08H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_NONE   | Retired load uops whose data source was LLC hit with no snoop required.   | Supports PEBS.  |
| D3H        | 01H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.LOCAL_DRAM | Retired load uops whose data source was local memory (cross-socket snoop not needed or missed).   | Supports PEBS.  |
| E6H        | 1FH         | BACLEARS.ANY                              | Number of front end re-steers due to BPU misprediction.   |   |

**Table 19-13. Performance Events In the Processor Core of 3rd Generation Intel® Core™ i7, i5, i3 Processors (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description   | Comment                          |
|------------|-------------|---------------------------|---|----------------------------------|
| F0H        | 01H         | L2_TRANS.DEMAND_DATA_RD   | Demand Data Read requests that access L2 cache.             |                                  |
| F0H        | 02H         | L2_TRANS.RFO              | RFO requests that access L2 cache.                          |                                  |
| F0H        | 04H         | L2_TRANS.CODE_RD          | L2 cache accesses when fetching instructions.               |                                  |
| F0H        | 08H         | L2_TRANS.ALL_PF           | Any MLC or LLC HW prefetch accessing L2, including rejects. |                                  |
| F0H        | 10H         | L2_TRANS.L1D_WB           | L1D writebacks that access L2 cache.                        |                                  |
| F0H        | 20H         | L2_TRANS.L2_FILL          | L2 fill requests that access L2 cache.                      |                                  |
| F0H        | 40H         | L2_TRANS.L2_WB            | L2 writebacks that access L2 cache.                         |                                  |
| F0H        | 80H         | L2_TRANS.ALL_REQUESTS     | Transactions accessing L2 pipe.                             |                                  |
| F1H        | 01H         | L2_LINES_IN.I             | L2 cache lines in I state filling L2.                       | Counting does not cover rejects. |
| F1H        | 02H         | L2_LINES_IN.S             | L2 cache lines in S state filling L2.                       | Counting does not cover rejects. |
| F1H        | 04H         | L2_LINES_IN.E             | L2 cache lines in E state filling L2.                       | Counting does not cover rejects. |
| F1H        | 07H         | L2_LINES_IN.ALL           | L2 cache lines filling L2.                                  | Counting does not cover rejects. |
| F2H        | 01H         | L2_LINES_OUT.DEMAND_CLEAN | Clean L2 cache lines evicted by demand.                     |                                  |
| F2H        | 02H         | L2_LINES_OUT.DEMAND_DIRTY | Dirty L2 cache lines evicted by demand.                     |                                  |
| F2H        | 04H         | L2_LINES_OUT.PF_CLEAN     | Clean L2 cache lines evicted by the MLC prefetcher.         |                                  |
| F2H        | 08H         | L2_LINES_OUT.PF_DIRTY     | Dirty L2 cache lines evicted by the MLC prefetcher.         |                                  |
| F2H        | 0AH         | L2_LINES_OUT.DIRTY_ALL    | Dirty L2 cache lines filling the L2.                        | Counting does not cover rejects. |

### 19.7.1 Performance Monitoring Events in the Processor Core of Intel Xeon Processor E5 v2 Family and Intel Xeon Processor E7 v2 Family

Model-specific performance monitoring events in the processor core that are applicable only to Intel Xeon processor E5 v2 family and Intel Xeon processor E7 v2 family based on the Ivy Bridge-E microarchitecture, with CPUID signature of DisplayFamily\_DisplayModel 06\_3EH, are listed in Table 19-14.

**Table 19-14. Performance Events Applicable Only to the Processor Core of Intel® Xeon® Processor E5 v2 Family and Intel® Xeon® Processor E7 v2 Family**

| Event Num. | Umask Value | Event Mask Mnemonic                        | Description   | Comment        |
|------------|-------------|--|---|----------------|
| D3H        | 03H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.LOCAL_DRAM  | Retired load uops whose data sources were local DRAM (snoop not needed, Snoop Miss, or Snoop Hit data not forwarded). | Supports PEBS. |
| D3H        | 0CH         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.REMOTE_DRAM | Retired load uops whose data source was remote DRAM (snoop not needed, Snoop Miss, or Snoop Hit data not forwarded).  | Supports PEBS. |
| D3H        | 10H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.REMOTE_HITM | Retired load uops whose data sources were remote HITM.  | Supports PEBS. |
| D3H        | 20H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.REMOTE_FWD  | Retired load uops whose data sources were forwards from a remote cache.   | Supports PEBS. |

## 19.8 PERFORMANCE MONITORING EVENTS FOR 2ND GENERATION INTEL® CORE™ I7-2XXX, INTEL® CORE™ I5-2XXX, INTEL® CORE™ I3-2XXX PROCESSOR SERIES

2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series, and Intel Xeon processor E3-1200 product family are based on the Intel microarchitecture code name Sandy Bridge. They support architectural performance monitoring events listed in Table 19-1. Model-specific performance monitoring events in the processor core are listed in Table 19-15, Table 19-16, and Table 19-17. The events in Table 19-15 apply to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_2AH and 06\_2DH. The events in Table 19-16 apply to processors with CPUID signature 06\_2AH. The events in Table 19-17 apply to processors with CPUID signature 06\_2DH. Fixed counters in the core PMU support the architecture events defined in Table 19-2.

Additional information on event specifics (e.g. derivative events using specific IA32\_PERFEVTSELx modifiers, limitations, special notes and recommendations) can be found at <https://software.intel.com/en-us/forums/software-tuning-performance-optimization-platform-monitoring>.

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description   | Comment                        |
|------------|-------------|-------------------------------------|---|--------------------------------|
| 03H        | 01H         | LD_BLOCKS.DATA_UNKNOWN              | Blocked loads due to store buffer blocks with unknown data.   |                                |
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD             | Loads blocked by overlapping with store buffer that cannot be forwarded.  |                                |
| 03H        | 08H         | LD_BLOCKS.NO_SR                     | # of Split loads blocked due to resource not available.   |                                |
| 03H        | 10H         | LD_BLOCKS.ALL_BLOCK                 | Number of cases where any load is blocked but has no DCU miss.  |                                |
| 05H        | 01H         | MISALIGN_MEM_REF.LOADS              | Speculative cache-line split load uops dispatched to L1D.   |                                |
| 05H        | 02H         | MISALIGN_MEM_REF.STORES             | Speculative cache-line split Store-address uops dispatched to L1D.  |                                |
| 07H        | 01H         | LD_BLOCKS_PARTIAL.ADDRESS_ALIAS     | False dependencies in MOB due to partial compare on address.  |                                |
| 07H        | 08H         | LD_BLOCKS_PARTIAL.ALL_STORE_BLOCK   | The number of times that load operations are temporarily blocked because of older stores, with addresses that are not yet known. A load operation may incur more than one block of this type. |                                |
| 08H        | 01H         | DTLB_LOAD_MISSES.MISS_CAUSES_A_WALK | Misses in all TLB levels that cause a page walk of any page size.   |                                |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED     | Misses in all TLB levels that caused page walk completed of any size.   |                                |
| 08H        | 04H         | DTLB_LOAD_MISSES.WALK_DURATION      | Cycle PMH is busy with a walk.  |                                |
| 08H        | 10H         | DTLB_LOAD_MISSES.STLB_HIT           | Number of cache load STLB hits. No page walk.   |                                |
| 0DH        | 03H         | INT_MISC.RECOVERY_CYCLES            | Cycles waiting to recover after Machine Clears or JEClear. Set Cmask= 1.  | Set Edge to count occurrences. |
| 0DH        | 40H         | INT_MISC.RAT_STALL_CYCLES           | Cycles RAT external stall is sent to IDQ for this thread.   |                                |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description   | Comment   |
|------------|-------------|--------------------------------------|---|---|
| 0EH        | 01H         | UOPS_ISSUED.ANY                      | Increments each cycle the # of Uops issued by the RAT to RS. Set Cmask = 1, Inv = 1, Any= 1 to count stalled cycles of this core. | Set Cmask = 1, Inv = 1 to count stalled cycles. |
| 10H        | 01H         | FP_COMP_OPS_EXE.X87                  | Counts number of X87 uops executed.   |   |
| 10H        | 10H         | FP_COMP_OPS_EXE.SSE_FP_PACKED_DOUBLE | Counts number of SSE* double precision FP packed uops executed.   |   |
| 10H        | 20H         | FP_COMP_OPS_EXE.SSE_FP_SCALAR_SINGLE | Counts number of SSE* single precision FP scalar uops executed.   |   |
| 10H        | 40H         | FP_COMP_OPS_EXE.SSE_PACKED_SINGLE    | Counts number of SSE* single precision FP packed uops executed.   |   |
| 10H        | 80H         | FP_COMP_OPS_EXE.SSE_SCALAR_DOUBLE    | Counts number of SSE* double precision FP scalar uops executed.   |   |
| 11H        | 01H         | SIMD_FP_256.PACKED_SINGLE            | Counts 256-bit packed single-precision floating-point instructions.   |   |
| 11H        | 02H         | SIMD_FP_256.PACKED_DOUBLE            | Counts 256-bit packed double-precision floating-point instructions.   |   |
| 14H        | 01H         | ARITH.FPU_DIV_ACTIVE                 | Cycles that the divider is active, includes INT and FP. Set 'edge =1, cmask=1' to count the number of divides.                    |   |
| 17H        | 01H         | INSTS_WRITTEN_TO_IQ.INSTS            | Counts the number of instructions written into the IQ every cycle.  |   |
| 24H        | 01H         | L2_RQSTS.DEMAND_DATA_READ_HIT        | Demand Data Read requests that hit L2 cache.  |   |
| 24H        | 03H         | L2_RQSTS.ALL_DEMAND_DATA_READ        | Counts any demand and L1 HW prefetch data load requests to L2.  |   |
| 24H        | 04H         | L2_RQSTS.RFO_HITS                    | Counts the number of store RFO requests that hit the L2 cache.  |   |
| 24H        | 08H         | L2_RQSTS.RFO_MISS                    | Counts the number of store RFO requests that miss the L2 cache.   |   |
| 24H        | 0CH         | L2_RQSTS.ALL_RFO                     | Counts all L2 store RFO requests.   |   |
| 24H        | 10H         | L2_RQSTS.CODE_READ_HIT               | Number of instruction fetches that hit the L2 cache.  |   |
| 24H        | 20H         | L2_RQSTS.CODE_READ_MISS              | Number of instruction fetches that missed the L2 cache.   |   |
| 24H        | 30H         | L2_RQSTS.ALL_CODE_READ               | Counts all L2 code requests.  |   |
| 24H        | 40H         | L2_RQSTS.PF_HIT                      | Requests from L2 Hardware prefetcher that hit L2.   |   |
| 24H        | 80H         | L2_RQSTS.PF_MISS                     | Requests from L2 Hardware prefetcher that missed L2.  |   |
| 24H        | C0H         | L2_RQSTS.ALL_PF                      | Any requests from L2 Hardware prefetchers.  |   |
| 27H        | 01H         | L2_STORE_LOCK_RQSTS.MISS             | RF0s that miss cache lines.   |   |
| 27H        | 04H         | L2_STORE_LOCK_RQSTS.HIT_E            | RF0s that hit cache lines in E state.   |   |
| 27H        | 08H         | L2_STORE_LOCK_RQSTS.HIT_M            | RF0s that hit cache lines in M state.   |   |
| 27H        | 0FH         | L2_STORE_LOCK_RQSTS.ALL              | RF0s that access cache lines in any state.  |   |



**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description  | Comment   |
|------------|-------------|--------------------------------------|--|---|
| 28H        | 01H         | L2_L1D_WB_RQSTS.MISS                 | Not rejected writebacks from L1D to L2 cache lines that missed L2.   |   |
| 28H        | 02H         | L2_L1D_WB_RQSTS.HIT_S                | Not rejected writebacks from L1D to L2 cache lines in S state.   |   |
| 28H        | 04H         | L2_L1D_WB_RQSTS.HIT_E                | Not rejected writebacks from L1D to L2 cache lines in E state.   |   |
| 28H        | 08H         | L2_L1D_WB_RQSTS.HIT_M                | Not rejected writebacks from L1D to L2 cache lines in M state.   |   |
| 28H        | 0FH         | L2_L1D_WB_RQSTS.ALL                  | Not rejected writebacks from L1D to L2 cache.  |   |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE          | This event counts requests originating from the core that reference a cache line in the last level cache.  | See Table 19-1.   |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS               | This event counts each cache miss condition for references to the last level cache.  | See Table 19-1.   |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P            | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. | See Table 19-1.   |
| 3CH        | 01H         | CPU_CLK_THREAD_UNHALTED.REF_XCLK     | Increments at the frequency of XCLK (100 MHz) when not halted.   | See Table 19-1.   |
| 48H        | 01H         | L1D_PEND_MISS.PENDING                | Increments the number of outstanding L1D misses every cycle. Set Cmask = 1 and Edge = 1 to count occurrences.  | PMC2 only;<br>Set Cmask = 1 to count cycles.                          |
| 49H        | 01H         | DTLB_STORE_MISSES.MISS_CAUSES_A_WALK | Miss in all TLB levels causes a page walk of any page size (4K/2M/4M/1G).  |   |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED     | Miss in all TLB levels causes a page walk that completes of any page size (4K/2M/4M/1G).   |   |
| 49H        | 04H         | DTLB_STORE_MISSES.WALK_DURATION      | Cycles PMH is busy with this walk.   |   |
| 49H        | 10H         | DTLB_STORE_MISSES.STLB_HIT           | Store operations that miss the first TLB level but hit the second and do not cause page walks.   |   |
| 4CH        | 01H         | LOAD_HIT_PRE.SW_PF                   | Not SW-prefetch load dispatches that hit fill buffer allocated for S/W prefetch.   |   |
| 4CH        | 02H         | LOAD_HIT_PRE.HW_PF                   | Not SW-prefetch load dispatches that hit fill buffer allocated for H/W prefetch.   |   |
| 4EH        | 02H         | HW_PRE_REQ.DL1_MISS                  | Hardware Prefetch requests that miss the L1D cache. A request is being counted each time it access the cache & miss it, including if a block is applicable or if hit the Fill Buffer for example.  | This accounts for both L1 streamer and IP-based (IPP) HW prefetchers. |
| 51H        | 01H         | L1D.REPLACEMENT                      | Counts the number of lines brought into the L1 data cache.   |   |
| 51H        | 02H         | L1D.ALLOCATED_IN_M                   | Counts the number of allocations of modified L1D cache lines.  |   |
| 51H        | 04H         | L1D.EVICTION                         | Counts the number of modified lines evicted from the L1 data cache due to replacement.   |   |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description   | Comment                        |
|------------|-------------|---|---|--------------------------------|
| 51H        | 08H         | L1D.ALL_M_REPLACEMENT                       | Cache lines in M state evicted out of L1D due to Snoop HitM or dirty line replacement.  |                                |
| 59H        | 20H         | PARTIAL_RAT_STALLS.FLAGS_MERGE_UOP          | Increments the number of flags-merge uops in flight each cycle. Set Cmask = 1 to count cycles.  |                                |
| 59H        | 40H         | PARTIAL_RAT_STALLS.SLOW_LEA_WINDOW          | Cycles with at least one slow LEA uop allocated.  |                                |
| 59H        | 80H         | PARTIAL_RAT_STALLS.MUL_SINGLE_UOP           | Number of Multiply packed/scalar single precision uops allocated.   |                                |
| 5BH        | 0CH         | RESOURCE_STALLS2.ALL_FL_EMPTY               | Cycles stalled due to free list empty.  | PMCO-3 only regardless HTT.    |
| 5BH        | 0FH         | RESOURCE_STALLS2.ALL_PRF_CONTROL            | Cycles stalled due to control structures full for physical registers.   |                                |
| 5BH        | 40H         | RESOURCE_STALLS2.BOB_FULL                   | Cycles Allocator is stalled due Branch Order Buffer.  |                                |
| 5BH        | 4FH         | RESOURCE_STALLS2.OOO_RESOURCE               | Cycles stalled due to out of order resources full.  |                                |
| 5CH        | 01H         | CPL_CYCLES.RING0                            | Unhalted core cycles when the thread is in ring 0.  | Use Edge to count transition.  |
| 5CH        | 02H         | CPL_CYCLES.RING123                          | Unhalted core cycles when the thread is not in ring 0.  |                                |
| 5EH        | 01H         | RS_EVENTS.EMPTY_CYCLES                      | Cycles the RS is empty for the thread.  |                                |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA_RD | Offcore outstanding Demand Data Read transactions in SQ to uncore. Set Cmask=1 to count cycles.   |                                |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND_RFO     | Offcore outstanding RFO store transactions in SQ to uncore. Set Cmask=1 to count cycles.  |                                |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ALL_DATA_RD    | Offcore outstanding cacheable data read transactions in SQ to uncore. Set Cmask=1 to count cycles.  |                                |
| 63H        | 01H         | LOCK_CYCLES.SPLIT_LOCK_UC_LOCK_DURATION     | Cycles in which the L1D and L2 are locked, due to a UC lock or split lock.  |                                |
| 63H        | 02H         | LOCK_CYCLES.CACHE_LOCK_DURATION             | Cycles in which the L1D is locked.  |                                |
| 79H        | 02H         | IDQ.EMPTY                                   | Counts cycles the IDQ is empty.   |                                |
| 79H        | 04H         | IDQ.MITE_UOPS                               | Increment each cycle # of uops delivered to IDQ from MITE path. Set Cmask = 1 to count cycles.  | Can combine Umask 04H and 20H. |
| 79H        | 08H         | IDQ.DSB_UOPS                                | Increment each cycle. # of uops delivered to IDQ from DSB path. Set Cmask = 1 to count cycles.  | Can combine Umask 08H and 10H. |
| 79H        | 10H         | IDQ.MS_DSB_UOPS                             | Increment each cycle # of uops delivered to IDQ when MS busy by DSB. Set Cmask = 1 to count cycles MS is busy. Set Cmask=1 and Edge =1 to count MS activations. | Can combine Umask 08H and 10H. |
| 79H        | 20H         | IDQ.MS_MITE_UOPS                            | Increment each cycle # of uops delivered to IDQ when MS is busy by MITE. Set Cmask = 1 to count cycles.   | Can combine Umask 04H and 20H. |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                           | Description   | Comment                             |
|------------|-------------|---|---|-------------------------------------|
| 79H        | 30H         | IDQ.MS_UOPS                                   | Increment each cycle # of uops delivered to IDQ from MS by either DSB or MITE. Set Cmask = 1 to count cycles. | Can combine Umask 04H, 08H and 30H. |
| 80H        | 02H         | ICACHE.MISSES                                 | Number of Instruction Cache, Streaming Buffer and Victim Cache Misses. Includes UC accesses.                  |                                     |
| 85H        | 01H         | ITLB_MISSES.MISS_CAUSES_A_WALK                | Misses in all ITLB levels that cause page walks.  |                                     |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETED                    | Misses in all ITLB levels that cause completed page walks.  |                                     |
| 85H        | 04H         | ITLB_MISSES.WALK_DURATION                     | Cycle PMH is busy with a walk.  |                                     |
| 85H        | 10H         | ITLB_MISSES.STLB_HIT                          | Number of cache load STLB hits. No page walk.   |                                     |
| 87H        | 01H         | ILD_STALL.LCP                                 | Stalls caused by changing prefix length of the instruction.   |                                     |
| 87H        | 04H         | ILD_STALL.IQ_FULL                             | Stall cycles due to IQ is full.   |                                     |
| 88H        | 41H         | BR_INST_EXEC.NONTAKEN_CONDITIONAL             | Not-taken macro conditional branches.   |                                     |
| 88H        | 81H         | BR_INST_EXEC.TAKEN_CONDITIONAL                | Taken speculative and retired conditional branches.   |                                     |
| 88H        | 82H         | BR_INST_EXEC.TAKEN_DIRECT_JUMP                | Taken speculative and retired conditional branches excluding calls and indirects.                             |                                     |
| 88H        | 84H         | BR_INST_EXEC.TAKEN_INDIRECT_JUMP_NON_CALL_RET | Taken speculative and retired indirect branches excluding calls and returns.                                  |                                     |
| 88H        | 88H         | BR_INST_EXEC.TAKEN_INDIRECT_NEAR_RETURN       | Taken speculative and retired indirect branches that are returns.   |                                     |
| 88H        | 90H         | BR_INST_EXEC.TAKEN_DIRECT_NEAR_CALL           | Taken speculative and retired direct near calls.  |                                     |
| 88H        | A0H         | BR_INST_EXEC.TAKEN_INDIRECT_NEAR_CALL         | Taken speculative and retired indirect near calls.  |                                     |
| 88H        | C1H         | BR_INST_EXEC.ALL_CONDITIONAL                  | Speculative and retired conditional branches.   |                                     |
| 88H        | C2H         | BR_INST_EXEC.ALL_DIRECT_JUMP                  | Speculative and retired conditional branches excluding calls and indirects.                                   |                                     |
| 88H        | C4H         | BR_INST_EXEC.ALL_INDIRECT_JUMP_NON_CALL_RET   | Speculative and retired indirect branches excluding calls and returns.  |                                     |
| 88H        | C8H         | BR_INST_EXEC.ALL_INDIRECT_NEAR_RETURN         | Speculative and retired indirect branches that are returns.   |                                     |
| 88H        | D0H         | BR_INST_EXEC.ALL_NEAR_CALL                    | Speculative and retired direct near calls.  |                                     |
| 88H        | FFH         | BR_INST_EXEC.ALL_BRANCHES                     | Speculative and retired branches.   |                                     |
| 89H        | 41H         | BR_MISP_EXEC.NONTAKEN_CONDITIONAL             | Not-taken mispredicted macro conditional branches.  |                                     |
| 89H        | 81H         | BR_MISP_EXEC.TAKEN_CONDITIONAL                | Taken speculative and retired mispredicted conditional branches.  |                                     |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                           | Description   | Comment                       |
|------------|-------------|---|---|-------------------------------|
| 89H        | 84H         | BR_MISP_EXEC.TAKEN_INDIRECT_JUMP_NON_CALL_RET | Taken speculative and retired mispredicted indirect branches excluding calls and returns.                                 |                               |
| 89H        | 88H         | BR_MISP_EXEC.TAKEN_RETURN_NEAR                | Taken speculative and retired mispredicted indirect branches that are returns.  |                               |
| 89H        | 90H         | BR_MISP_EXEC.TAKEN_DIRECT_NEAR_CALL           | Taken speculative and retired mispredicted direct near calls.   |                               |
| 89H        | A0H         | BR_MISP_EXEC.TAKEN_INDIRECT_NEAR_CALL         | Taken speculative and retired mispredicted indirect near calls.   |                               |
| 89H        | C1H         | BR_MISP_EXEC.ALL_CONDITIONAL                  | Speculative and retired mispredicted conditional branches.  |                               |
| 89H        | C4H         | BR_MISP_EXEC.ALL_INDIRECT_JUMP_NON_CALL_RET   | Speculative and retired mispredicted indirect branches excluding calls and returns.                                       |                               |
| 89H        | D0H         | BR_MISP_EXEC.ALL_NEAR_CALL                    | Speculative and retired mispredicted direct near calls.   |                               |
| 89H        | FFH         | BR_MISP_EXEC.ALL_BRANCHES                     | Speculative and retired mispredicted branches.  |                               |
| 9CH        | 01H         | IDQ_UOPS_NOT_DELIVERED.CORE                   | Count issue pipeline slots where no uop was delivered from the front end to the back end when there is no back-end stall. | Use Cmask to qualify uop b/w. |
| A1H        | 01H         | UOPS_DISPATCHED_PORT.PORT_0                   | Cycles which a Uop is dispatched on port 0.   |                               |
| A1H        | 02H         | UOPS_DISPATCHED_PORT.PORT_1                   | Cycles which a Uop is dispatched on port 1.   |                               |
| A1H        | 0CH         | UOPS_DISPATCHED_PORT.PORT_2                   | Cycles which a Uop is dispatched on port 2.   |                               |
| A1H        | 30H         | UOPS_DISPATCHED_PORT.PORT_3                   | Cycles which a Uop is dispatched on port 3.   |                               |
| A1H        | 40H         | UOPS_DISPATCHED_PORT.PORT_4                   | Cycles which a Uop is dispatched on port 4.   |                               |
| A1H        | 80H         | UOPS_DISPATCHED_PORT.PORT_5                   | Cycles which a Uop is dispatched on port 5.   |                               |
| A2H        | 01H         | RESOURCE_STALLS.ANY                           | Cycles Allocation is stalled due to Resource Related reason.  |                               |
| A2H        | 02H         | RESOURCE_STALLS.LB                            | Counts the cycles of stall due to lack of load buffers.   |                               |
| A2H        | 04H         | RESOURCE_STALLS.RS                            | Cycles stalled due to no eligible RS entry available.   |                               |
| A2H        | 08H         | RESOURCE_STALLS.SB                            | Cycles stalled due to no store buffers available (not including draining from sync).                                      |                               |
| A2H        | 10H         | RESOURCE_STALLS.ROB                           | Cycles stalled due to re-order buffer full.   |                               |
| A2H        | 20H         | RESOURCE_STALLS.FCSW                          | Cycles stalled due to writing the FPU control word.   |                               |
| A3H        | 01H         | CYCLE_ACTIVITY.CYCLES_L2_PENDING              | Cycles with pending L2 miss loads. Set AnyThread to count per core.   |                               |
| A3H        | 02H         | CYCLE_ACTIVITY.CYCLES_L1D_PENDING             | Cycles with pending L1 cache miss loads. Set AnyThread to count per core.   | PMC2 only.                    |
| A3H        | 04H         | CYCLE_ACTIVITY.CYCLES_NO_DISPATCH             | Cycles of dispatch stalls. Set AnyThread to count per core.   | PMCO-3 only.                  |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description   | Comment                             |
|------------|-------------|--|---|-------------------------------------|
| A3H        | 05H         | CYCLE_ACTIVITY.STALL_CYCLE_S_L2_PENDING  |   | PMCO-3 only.                        |
| A3H        | 06H         | CYCLE_ACTIVITY.STALL_CYCLE_S_L1D_PENDING |   | PMC2 only.                          |
| A8H        | 01H         | LSD.UOPS                                 | Number of Uops delivered by the LSD.  |                                     |
| ABH        | 01H         | DSB2MITE_SWITCHES.COUNT                  | Number of DSB to MITE switches.   |                                     |
| ABH        | 02H         | DSB2MITE_SWITCHES.PENALTY_CYCLES         | Cycles DSB to MITE switches caused delay.   |                                     |
| ACH        | 02H         | DSB_FILL.OTHER_CANCEL                    | Cases of cancelling valid DSB fill not because of exceeding way limit.  |                                     |
| ACH        | 08H         | DSB_FILL.EXCEED_DSB_LINES                | DSB Fill encountered > 3 DSB lines.   |                                     |
| AEH        | 01H         | ITLB.ITLB_FLUSH                          | Counts the number of ITLB flushes; includes 4k/2M/4M pages.   |                                     |
| BOH        | 01H         | OFFCORE_REQUESTS.DEMAND_DATA_RD          | Demand data read requests sent to uncore.   |                                     |
| BOH        | 04H         | OFFCORE_REQUESTS.DEMAND_RFO              | Demand RFO read requests sent to uncore, including regular RFOs, locks, ItoM.   |                                     |
| BOH        | 08H         | OFFCORE_REQUESTS.ALL_DATA_RD             | Data read requests sent to uncore (demand and prefetch).  |                                     |
| B1H        | 01H         | UOPS_DISPATCHED.THREAD                   | Counts total number of uops to be dispatched per-thread each cycle. Set Cmask = 1, INV =1 to count stall cycles.  | PMCO-3 only regardless HTT.         |
| B1H        | 02H         | UOPS_DISPATCHED.CORE                     | Counts total number of uops to be dispatched per-core each cycle.   | Do not need to set ANY.             |
| B2H        | 01H         | OFFCORE_REQUESTS_BUFFER_SQ_FULL          | Offcore requests buffer cannot take more entries for this thread core.  |                                     |
| B6H        | 01H         | AGU_BYPASS_CANCEL.COUNT                  | Counts executed load operations with all the following traits: 1. Addressing of the format [base + offset], 2. The offset is between 1 and 2047, 3. The address specified in the base register is in one page and the address [base+offset] is in another page. |                                     |
| B7H        | 01H         | OFF_CORE_RESPONSE_0                      | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A6H.                 |
| BBH        | 01H         | OFF_CORE_RESPONSE_1                      | See Section 18.3.4.5, "Off-core Response Performance Monitoring".   | Requires MSR 01A7H.                 |
| BDH        | 01H         | TLB_FLUSH.DTLB_THREAD                    | DTLB flush attempts of the thread-specific entries.   |                                     |
| BDH        | 20H         | TLB_FLUSH.STLB_ANY                       | Count number of STLB flush attempts.  |                                     |
| BFH        | 05H         | L1D_BLOCKS.BANK_CONFLICT_CYCLES          | Cycles when dispatched loads are cancelled due to L1D bank conflicts with other load ports.   | Cmask=1.                            |
| COH        | 00H         | INST_RETIRED.ANY_P                       | Number of instructions at retirement.   | See Table 19-1.                     |
| COH        | 01H         | INST_RETIRED.PREC_DIST                   | Precise instruction retired event with HW to reduce effect of PEBS shadow in IP distribution.   | PMC1 only; must quiesce other PMCs. |
| C1H        | 02H         | OTHER_ASSISTS.ITLB_MISS_RETIRED          | Instructions that experienced an ITLB miss.   |                                     |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment         |
|------------|-------------|--------------------------------|--|-----------------|
| C1H        | 08H         | OTHER_ASSISTS.AVX_STORE        | Number of assists associated with 256-bit AVX store operations.  |                 |
| C1H        | 10H         | OTHER_ASSISTS.AVX_TO_SSE       | Number of transitions from AVX-256 to legacy SSE when penalty applicable.  |                 |
| C1H        | 20H         | OTHER_ASSISTS.SSE_TO_AVX       | Number of transitions from SSE to AVX-256 when penalty applicable.   |                 |
| C2H        | 01H         | UOPS_RETIRED.ALL               | Counts the number of micro-ops retired, Use cmask=1 and invert to count active cycles or stalled cycles.                     | Supports PEBS.  |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS      | Counts the number of retirement slots used each cycle.   | Supports PEBS.  |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING | Counts the number of machine clears due to memory order conflicts.   |                 |
| C3H        | 04H         | MACHINE_CLEARS.SMC             | Counts the number of times that a program writes to a code section.  |                 |
| C3H        | 20H         | MACHINE_CLEARS.MASKMOV         | Counts the number of executed AVX masked load operations that refer to an illegal address range with the mask bits set to 0. |                 |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES   | Branch instructions at retirement.   | See Table 19-1. |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL    | Counts the number of conditional branch instructions retired.  | Supports PEBS.  |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL      | Direct and indirect near call instructions retired.  | Supports PEBS.  |
| C4H        | 04H         | BR_INST_RETIRED.ALL_BRANCHES   | Counts the number of branch instructions retired.  | Supports PEBS.  |
| C4H        | 08H         | BR_INST_RETIRED.NEAR_RETURN    | Counts the number of near return instructions retired.   | Supports PEBS.  |
| C4H        | 10H         | BR_INST_RETIRED.NOT_TAKEN      | Counts the number of not taken branch instructions retired.  |                 |
| C4H        | 20H         | BR_INST_RETIRED.NEAR_TAKEN     | Number of near taken branches retired.   | Supports PEBS.  |
| C4H        | 40H         | BR_INST_RETIRED.FAR_BRANCH     | Number of far branches retired.  |                 |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted branch instructions at retirement.  | See Table 19-1. |
| C5H        | 01H         | BR_MISP_RETIRED.CONDITIONAL    | Mispredicted conditional branch instructions retired.  | Supports PEBS.  |
| C5H        | 02H         | BR_MISP_RETIRED.NEAR_CALL      | Direct and indirect mispredicted near call instructions retired.   | Supports PEBS.  |
| C5H        | 04H         | BR_MISP_RETIRED.ALL_BRANCHES   | Mispredicted macro branch instructions retired.  | Supports PEBS.  |
| C5H        | 10H         | BR_MISP_RETIRED.NOT_TAKEN      | Mispredicted not taken branch instructions retired.  | Supports PEBS.  |
| C5H        | 20H         | BR_MISP_RETIRED.TAKEN          | Mispredicted taken branch instructions retired.  | Supports PEBS.  |
| CAH        | 02H         | FP_ASSIST.X87_OUTPUT           | Number of X87 assists due to output value.   |                 |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description  | Comment                                    |
|------------|-------------|---|--|--|
| CAH        | 04H         | FP_ASSIST.X87_INPUT                         | Number of X87 assists due to input value.  |  |
| CAH        | 08H         | FP_ASSIST.SIMD_OUTPUT                       | Number of SIMD FP assists due to output values.  |  |
| CAH        | 10H         | FP_ASSIST.SIMD_INPUT                        | Number of SIMD FP assists due to input values.   |  |
| CAH        | 1EH         | FP_ASSIST.ANY                               | Cycles with any input/output SSE* or FP assists.   |  |
| CCH        | 20H         | ROB_MISC_EVENTS.LBR_INSE<br>RTS             | Count cases of saving new LBR records by hardware.   |  |
| CDH        | 01H         | MEM_TRANS_RETIRED.LOAD_<br>LATENCY          | Randomly sampled loads whose latency is above a user defined threshold. A small fraction of the overall loads are sampled due to randomization. PMC3 only. | Specify threshold in MSR 3F6H.             |
| CDH        | 02H         | MEM_TRANS_RETIRED.PRECIS<br>E_STORE         | Sample stores and collect precise store operation via PEBS record. PMC3 only.  | See Section 18.3.4.4.3.                    |
| DOH        | 11H         | MEM_UOPS_RETIRED.STLB_MI<br>SS_LOADS        | Retired load uops that miss the STLB.  | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 12H         | MEM_UOPS_RETIRED.STLB_MI<br>SS_STORES       | Retired store uops that miss the STLB.   | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 21H         | MEM_UOPS_RETIRED.LOCK_LO<br>ADS             | Retired load uops with locked access.  | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 41H         | MEM_UOPS_RETIRED.SPLIT_L<br>OADS            | Retired load uops that split across a cacheline boundary.  | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 42H         | MEM_UOPS_RETIRED.SPLIT_S<br>TORES           | Retired store uops that split across a cacheline boundary.   | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 81H         | MEM_UOPS_RETIRED.ALL_LOA<br>DS              | All retired load uops.   | Supports PEBS. PMC0-3 only regardless HTT. |
| DOH        | 82H         | MEM_UOPS_RETIRED.ALL_STO<br>RES             | All retired store uops.  | Supports PEBS. PMC0-3 only regardless HTT. |
| D1H        | 01H         | MEM_LOAD_UOPS_RETIRED.L<br>1_HIT            | Retired load uops with L1 cache hits as data sources.  | Supports PEBS. PMC0-3 only regardless HTT. |
| D1H        | 02H         | MEM_LOAD_UOPS_RETIRED.L<br>2_HIT            | Retired load uops with L2 cache hits as data sources.  | Supports PEBS.                             |
| D1H        | 04H         | MEM_LOAD_UOPS_RETIRED.LL<br>C_HIT           | Retired load uops which data sources were data hits in LLC without snoops required.  | Supports PEBS.                             |
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRED.LL<br>C_MISS          | Retired load uops which data sources were data missed LLC (excluding unknown data source).   | Supports PEBS.                             |
| D1H        | 40H         | MEM_LOAD_UOPS_RETIRED.HI<br>T_LFB           | Retired load uops which data sources were load uops missed L1 but hit FB due to preceding miss to the same cache line with data not ready.                 | Supports PEBS.                             |
| D2H        | 01H         | MEM_LOAD_UOPS_LLC_HIT_R<br>ETIRED.XSNP_MISS | Retired load uops whose data source was an on-package core cache LLC hit and cross-core snoop missed.  | Supports PEBS.                             |
| D2H        | 02H         | MEM_LOAD_UOPS_LLC_HIT_R<br>ETIRED.XSNP_HIT  | Retired load uops whose data source was an on-package LLC hit and cross-core snoop hits.   | Supports PEBS.                             |
| D2H        | 04H         | MEM_LOAD_UOPS_LLC_HIT_R<br>ETIRED.XSNP_HITM | Retired load uops whose data source was an on-package core cache with HitM responses.  | Supports PEBS.                             |
| D2H        | 08H         | MEM_LOAD_UOPS_LLC_HIT_R<br>ETIRED.XSNP_NONE | Retired load uops whose data source was LLC hit with no snoop required.  | Supports PEBS.                             |

**Table 19-15. Performance Events In the Processor Core Common to 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series and Intel® Xeon® Processors E3 and E5 Family (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description   | Comment                          |
|------------|-------------|---------------------------|---|----------------------------------|
| E6H        | 01H         | BACLEAR.S.ANY             | Counts the number of times the front end is re-steered, mainly when the BPU cannot provide a correct prediction and this is corrected by other branch handling mechanisms at the front end. |                                  |
| F0H        | 01H         | L2_TRANS.DEMAND_DATA_RD   | Demand Data Read requests that access L2 cache.   |                                  |
| F0H        | 02H         | L2_TRANS.RFO              | RFO requests that access L2 cache.  |                                  |
| F0H        | 04H         | L2_TRANS.CODE_RD          | L2 cache accesses when fetching instructions.   |                                  |
| F0H        | 08H         | L2_TRANS.ALL_PF           | L2 or LLC HW prefetches that access L2 cache.   | Including rejects.               |
| F0H        | 10H         | L2_TRANS.L1D_WB           | L1D writebacks that access L2 cache.  |                                  |
| F0H        | 20H         | L2_TRANS.L2_FILL          | L2 fill requests that access L2 cache.  |                                  |
| F0H        | 40H         | L2_TRANS.L2_WB            | L2 writebacks that access L2 cache.   |                                  |
| F0H        | 80H         | L2_TRANS.ALL_REQUESTS     | Transactions accessing L2 pipe.   |                                  |
| F1H        | 01H         | L2_LINES_IN.I             | L2 cache lines in I state filling L2.   | Counting does not cover rejects. |
| F1H        | 02H         | L2_LINES_IN.S             | L2 cache lines in S state filling L2.   | Counting does not cover rejects. |
| F1H        | 04H         | L2_LINES_IN.E             | L2 cache lines in E state filling L2.   | Counting does not cover rejects. |
| F1H        | 07H         | L2_LINES_IN.ALL           | L2 cache lines filling L2.  | Counting does not cover rejects. |
| F2H        | 01H         | L2_LINES_OUT.DEMAND_CLEAN | Clean L2 cache lines evicted by demand.   |                                  |
| F2H        | 02H         | L2_LINES_OUT.DEMAND_DIRTY | Dirty L2 cache lines evicted by demand.   |                                  |
| F2H        | 04H         | L2_LINES_OUT.PF_CLEAN     | Clean L2 cache lines evicted by L2 prefetch.  |                                  |
| F2H        | 08H         | L2_LINES_OUT.PF_DIRTY     | Dirty L2 cache lines evicted by L2 prefetch.  |                                  |
| F2H        | 0AH         | L2_LINES_OUT.DIRTY_ALL    | Dirty L2 cache lines filling the L2.  | Counting does not cover rejects. |
| F4H        | 10H         | SQ_MISC.SPLIT_LOCK        | Split locks in SQ.  |                                  |

Non-architecture performance monitoring events in the processor core that are applicable only to Intel processors with CPUID signature of DisplayFamily\_DisplayModel 06\_2AH are listed in Table 19-16.

**Table 19-16. Performance Events applicable only to the Processor core for 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment                                    |
|------------|-------------|---|---|--|
| D2H        | 01H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_MISS | Retired load uops which data sources were LLC hit and cross-core snoop missed in on-pkg core cache. | Supports PEBS. PMCO-3 only regardless HTT. |
| D2H        | 02H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HIT  | Retired load uops which data sources were LLC and cross-core snoop hits in on-pkg core cache.       | Supports PEBS.                             |
| D2H        | 04H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM | Retired load uops which data sources were HitM responses from shared LLC.                           | Supports PEBS.                             |
| D2H        | 08H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_NONE | Retired load uops which data sources were hits in LLC without snoops required.                      | Supports PEBS.                             |



**Table 19-16. Performance Events applicable only to the Processor core for 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description  | Comment                                    |
|------------|-------------|---|--|--|
| D4H        | 02H         | MEM_LOAD_UOPS_MISC_RETI<br>RED.LLC_MISS                         | Retired load uops with unknown information as data source in cache serviced the load.  | Supports PEBS. PMCO-3 only regardless HTT. |
| B7H/BBH    | 01H         | OFFCORE_RESPONSE_N  | Sub-events of OFFCORE_RESPONSE_N (suffix N = 0, 1) programmed using MSR 01A6H/01A7H with values shown in the comment column. |  |
|            |             | OFFCORE_RESPONSE.ALL_CODE_RD.LLC_HIT_N                          |  | 10003C0244H                                |
|            |             | OFFCORE_RESPONSE.ALL_CODE_RD.LLC_HIT.NO_SNOOP_NEEDED_N          |  | 1003C0244H                                 |
|            |             | OFFCORE_RESPONSE.ALL_CODE_RD.LLC_HIT.SNOOP_MISS_N               |  | 2003C0244H                                 |
|            |             | OFFCORE_RESPONSE.ALL_CODE_RD.LLC_HIT.MISS_DRAM_N                |  | 300400244H                                 |
|            |             | OFFCORE_RESPONSE.ALL_DATA_RD.LLC_HIT.ANY_RESPONSE_N             |  | 3F803C0091H                                |
|            |             | OFFCORE_RESPONSE.ALL_DATA_RD.LLC_MISS.DRAM_N                    |  | 300400091H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_HIT.ANY_RESPONSE_N          |  | 3F803C0240H                                |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N |  | 4003C0240H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_HIT.HITM_OTHER_CORE_N       |  | 10003C0240H                                |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_HIT.NO_SNOOP_NEEDED_N       |  | 1003C0240H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_HIT.SNOOP_MISS_N            |  | 2003C0240H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_CODE_RD.LLC_MISS.DRAM_N                 |  | 300400240H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_DATA_RD.LLC_MISS.DRAM_N                 |  | 300400090H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_HIT.ANY_RESPONSE_N              |  | 3F803C0120H                                |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N     |  | 4003C0120H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_HIT.HITM_OTHER_CORE_N           |  | 10003C0120H                                |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_HIT.NO_SNOOP_NEEDED_N           |  | 1003C0120H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_HIT.SNOOP_MISS_N                |  | 2003C0120H                                 |
|            |             | OFFCORE_RESPONSE.ALL_PF_RFO.LLC_MISS.DRAM_N                     |  | 300400120H                                 |
|            |             | OFFCORE_RESPONSE.ALL_READS.LLC_MISS.DRAM_N                      |  | 3004003F7H                                 |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_HIT.ANY_RESPONSE_N                 |  | 3F803C0122H                                |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N        |  | 4003C0122H                                 |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_HIT.HITM_OTHER_CORE_N              |  | 10003C0122H                                |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_HIT.NO_SNOOP_NEEDED_N              |  | 1003C0122H                                 |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_HIT.SNOOP_MISS_N                   |  | 2003C0122H                                 |
|            |             | OFFCORE_RESPONSE.ALL_RFO.LLC_MISS.DRAM_N                        |  | 300400122H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N |  | 4003C0004H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_HIT.HITM_OTHER_CORE_N       |  | 10003C0004H                                |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_HIT.NO_SNOOP_NEEDED_N       |  | 1003C0004H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_HIT.SNOOP_MISS_N            |  | 2003C0004H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.DRAM_N                 |  | 300400004H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.DRAM_N                 |  | 300400001H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_HIT.ANY_RESPONSE_N              |  | 3F803C0002H                                |
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N     |  | 4003C0002H                                 |
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_HIT.HITM_OTHER_CORE_N           |  | 10003C0002H                                |

**Table 19-16. Performance Events applicable only to the Processor core for 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic   | Description | Comment     |
|------------|-------------|---|-------------|-------------|
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_HIT.NO_SNOOP_NEEDED_N           |             | 1003C0002H  |
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_HIT.SNOOP_MISS_N                |             | 2003C0002H  |
|            |             | OFFCORE_RESPONSE.DEMAND_RFO.LLC_MISS.DRAM_N                     |             | 300400002H  |
|            |             | OFFCORE_RESPONSE.OTHER.ANY_RESPONSE_N                           |             | 18000H      |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N  |             | 4003C0040H  |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_HIT.HITM_OTHER_CORE_N        |             | 10003C0040H |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_HIT.NO_SNOOP_NEEDED_N        |             | 1003C0040H  |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_HIT.SNOOP_MISS_N             |             | 2003C0040H  |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_MISS.DRAM_N                  |             | 300400040H  |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.DRAM_N                  |             | 300400010H  |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_HIT.ANY_RESPONSE_N               |             | 3F803C0020H |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N      |             | 4003C0020H  |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_HIT.HITM_OTHER_CORE_N            |             | 10003C0020H |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_HIT.NO_SNOOP_NEEDED_N            |             | 1003C0020H  |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_HIT.SNOOP_MISS_N                 |             | 2003C0020H  |
|            |             | OFFCORE_RESPONSE.PF_L2_RFO.LLC_MISS.DRAM_N                      |             | 300400020H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N |             | 4003C0200H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_HIT.HITM_OTHER_CORE_N       |             | 10003C0200H |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_HIT.NO_SNOOP_NEEDED_N       |             | 1003C0200H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_HIT.SNOOP_MISS_N            |             | 2003C0200H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_MISS.DRAM_N                 |             | 300400200H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_DATA_RD.LLC_MISS.DRAM_N                 |             | 300400080H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_HIT.ANY_RESPONSE_N              |             | 3F803C0100H |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_HIT.HIT_OTHER_CORE_NO_FWD_N     |             | 4003C0100H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_HIT.HITM_OTHER_CORE_N           |             | 10003C0100H |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_HIT.NO_SNOOP_NEEDED_N           |             | 1003C0100H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_HIT.SNOOP_MISS_N                |             | 2003C0100H  |
|            |             | OFFCORE_RESPONSE.PF_LLC_RFO.LLC_MISS.DRAM_N                     |             | 300400100H  |

Non-architecture performance monitoring events in the processor core that are applicable only to Intel Xeon processor E5 family (and Intel Core i7-3930 processor) based on Intel microarchitecture code name Sandy Bridge, with CPUID signature of DisplayFamily\_DisplayModel 06\_2DH, are listed in Table 19-17.

**Table 19-17. Performance Events Applicable only to the Processor Core of Intel® Xeon® Processor E5 Family**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment |
|------------|-------------|--------------------------------|---|---------|
| CDH        | 01H         | MEM_TRANS_RETIRED.LOAD_LATENCY | Additional Configuration: Disable BL bypass and direct2core, and if the memory is remotely homed. The count is not reliable If the memory is locally homed. |         |
| D1H        | 04H         | MEM_LOAD_UOPS_RETIRED.LLC_HIT  | Additional Configuration: Disable BL bypass. Supports PEBS.   |         |

**Table 19-17. Performance Events Applicable only to the Processor Core of Intel® Xeon® Processor E5 Family**

| Event Num. | Umask Value | Event Mask Mnemonic                                       | Description   | Comment   |
|------------|-------------|---|---|---|
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRED.LLC_MISS                            | Additional Configuration: Disable BL bypass and direct2core. Supports PEBS.   |   |
| D2H        | 01H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_MISS                   | Additional Configuration: Disable bypass. Supports PEBS.  |   |
| D2H        | 02H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HIT                    | Additional Configuration: Disable bypass. Supports PEBS.  |   |
| D2H        | 04H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM                   | Additional Configuration: Disable bypass. Supports PEBS.  |   |
| D2H        | 08H         | MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_NONE                   | Additional Configuration: Disable bypass. Supports PEBS.  |   |
| D3H        | 01H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.LOCAL_DRAM                 | Retired load uops which data sources were data missed LLC but serviced by local DRAM. Supports PEBS.                          | Disable BL bypass and direct2core (see MSR 3C9H). |
| D3H        | 04H         | MEM_LOAD_UOPS_LLC_MISS_RETIRED.REMOTE_DRAM                | Retired load uops which data sources were data missed LLC but serviced by remote DRAM. Supports PEBS.                         | Disable BL bypass and direct2core (see MSR 3C9H). |
| B7H/BBH    | 01H         | OFF_CORE_RESPONSE_N                                       | Sub-events of OFF_CORE_RESPONSE_N (suffix N = 0, 1) programmed using MSR 01A6H/01A7H with values shown in the comment column. |   |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.ANY_RESPONSE_N   |   | 3FFFC00004H                                       |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.LOCAL_DRAM_N     |   | 600400004H  |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.REMOTE_DRAM_N    |   | 67F800004H  |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.REMOTE_HIT_FWD_N |   | 87F800004H  |
|            |             | OFFCORE_RESPONSE.DEMAND_CODE_RD.LLC_MISS.REMOTE_HITM_N    |   | 107FC00004H                                       |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.ANY_DRAM_N       |   | 67FC00001H  |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.ANY_RESPONSE_N   |   | 3F803C0001H                                       |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.LOCAL_DRAM_N     |   | 600400001H  |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.REMOTE_DRAM_N    |   | 67F800001H  |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.REMOTE_HIT_FWD_N |   | 87F800001H  |
|            |             | OFFCORE_RESPONSE.DEMAND_DATA_RD.LLC_MISS.REMOTE_HITM_N    |   | 107FC00001H                                       |
|            |             | OFFCORE_RESPONSE.PF_L2_CODE_RD.LLC_MISS.ANY_RESPONSE_N    |   | 3F803C0040H                                       |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.ANY_DRAM_N        |   | 67FC00010H  |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.ANY_RESPONSE_N    |   | 3F803C0010H                                       |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.LOCAL_DRAM_N      |   | 600400010H  |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.REMOTE_DRAM_N     |   | 67F800010H  |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.REMOTE_HIT_FWD_N  |   | 87F800010H  |
|            |             | OFFCORE_RESPONSE.PF_L2_DATA_RD.LLC_MISS.REMOTE_HITM_N     |   | 107FC00010H                                       |
|            |             | OFFCORE_RESPONSE.PF_LLC_CODE_RD.LLC_MISS.ANY_RESPONSE_N   |   | 3FFFC00200H                                       |
|            |             | OFFCORE_RESPONSE.PF_LLC_DATA_RD.LLC_MISS.ANY_RESPONSE_N   |   | 3FFFC00080H                                       |

Model-specific performance monitoring events that are located in the uncore sub-system are implementation specific between different platforms using processors based on Intel microarchitecture code name Sandy Bridge. Processors with CPUID signature of DisplayFamily\_DisplayModel 06\_2AH support performance events listed in Table 19-18.

**Table 19-18. Performance Events In the Processor Uncore for 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series**

| Event Num. <sup>1</sup> | Umask Value | Event Mask Mnemonic                     | Description   | Comment  |
|-------------------------|-------------|---|---|--|
| 22H                     | 01H         | UNC_CBO_XSNP_RESPONSE.MISS              | A snoop misses in some processor core.  | Must combine with one of the umask values of 20H, 40H, 80H.      |
| 22H                     | 02H         | UNC_CBO_XSNP_RESPONSE.INVAL             | A snoop invalidates a non-modified line in some processor core.   |  |
| 22H                     | 04H         | UNC_CBO_XSNP_RESPONSE.HIT               | A snoop hits a non-modified line in some processor core.  |  |
| 22H                     | 08H         | UNC_CBO_XSNP_RESPONSE.HITM              | A snoop hits a modified line in some processor core.  |  |
| 22H                     | 10H         | UNC_CBO_XSNP_RESPONSE.INVAL_M           | A snoop invalidates a modified line in some processor core.   |  |
| 22H                     | 20H         | UNC_CBO_XSNP_RESPONSE.EXTERNAL_FILTER   | Filter on cross-core snoops initiated by this Cbox due to external snoop request.   | Must combine with at least one of 01H, 02H, 04H, 08H, 10H.       |
| 22H                     | 40H         | UNC_CBO_XSNP_RESPONSE.CORE_FILTER       | Filter on cross-core snoops initiated by this Cbox due to processor core memory request.  |  |
| 22H                     | 80H         | UNC_CBO_XSNP_RESPONSE.EVICTION_FILTER   | Filter on cross-core snoops initiated by this Cbox due to LLC eviction.   |  |
| 34H                     | 01H         | UNC_CBO_CACHE_LOOKUP.M                  | LLC lookup request that access cache and found line in M-state.   | Must combine with one of the umask values of 10H, 20H, 40H, 80H. |
| 34H                     | 02H         | UNC_CBO_CACHE_LOOKUP.E                  | LLC lookup request that access cache and found line in E-state.   |  |
| 34H                     | 04H         | UNC_CBO_CACHE_LOOKUP.S                  | LLC lookup request that access cache and found line in S-state.   |  |
| 34H                     | 08H         | UNC_CBO_CACHE_LOOKUP.I                  | LLC lookup request that access cache and found line in I-state.   |  |
| 34H                     | 10H         | UNC_CBO_CACHE_LOOKUP.READ_FILTER        | Filter on processor core initiated cacheable read requests. Must combine with at least one of 01H, 02H, 04H, 08H.   |  |
| 34H                     | 20H         | UNC_CBO_CACHE_LOOKUP.WRITE_FILTER       | Filter on processor core initiated cacheable write requests. Must combine with at least one of 01H, 02H, 04H, 08H.  |  |
| 34H                     | 40H         | UNC_CBO_CACHE_LOOKUP.EXTSNP_FILTER      | Filter on external snoop requests. Must combine with at least one of 01H, 02H, 04H, 08H.  |  |
| 34H                     | 80H         | UNC_CBO_CACHE_LOOKUP.ANY_REQUEST_FILTER | Filter on any IRQ or IPQ initiated requests including uncacheable, non-coherent requests. Must combine with at least one of 01H, 02H, 04H, 08H.   |  |
| 80H                     | 01H         | UNC_ARB_TRK_OCCUPANCY.ALL               | Counts cycles weighted by the number of requests waiting for data returning from the memory controller. Accounts for coherent and non-coherent requests initiated by IA cores, processor graphic units, or LLC. | Counter 0 only.  |
| 81H                     | 01H         | UNC_ARB_TRK_REQUEST.ALL                 | Counts the number of coherent and in-coherent requests initiated by IA cores, processor graphic units, or LLC.  |  |
| 81H                     | 20H         | UNC_ARB_TRK_REQUEST.WRITES              | Counts the number of allocated write entries, include full, partial, and LLC evictions.   |  |
| 81H                     | 80H         | UNC_ARB_TRK_REQUEST.EVICTIONS           | Counts the number of LLC evictions allocated.   |  |

**Table 19-18. Performance Events In the Processor Uncore for 2nd Generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx Processor Series (Contd.)**

| Event Num. <sup>1</sup> | Umask Value | Event Mask Mnemonic           | Description   | Comment         |
|-------------------------|-------------|-------------------------------|---|-----------------|
| 83H                     | 01H         | UNC_ARB_COH_TRK_OCCUPANCY.ALL | Cycles weighted by number of requests pending in Coherency Tracker. | Counter 0 only. |
| 84H                     | 01H         | UNC_ARB_COH_TRK_REQUESTS.ALL  | Number of requests allocated in Coherency Tracker.                  |                 |

**NOTES:**

1. The uncore events must be programmed using MSRs located in specific performance monitoring units in the uncore. UNC\_CBO\* events are supported using MSR\_UNC\_CBO\* MSRs; UNC\_ARB\* events are supported using MSR\_UNC\_ARB\*MSRs.

## 19.9 PERFORMANCE MONITORING EVENTS FOR INTEL® CORE™ I7 PROCESSOR FAMILY AND INTEL® XEON® PROCESSOR FAMILY

Processors based on the Intel microarchitecture code name Nehalem support the architectural and model-specific performance monitoring events listed in Table 19-1 and Table 19-19. The events in Table 19-19 generally applies to processors with CPUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_1AH, 06\_1EH, 06\_1FH, and 06\_2EH. However, Intel Xeon processors with CPUID signature of DisplayFamily\_DisplayModel 06\_2EH have a small number of events that are not supported in processors with CPUID signature 06\_1AH, 06\_1EH, and 06\_1FH. These events are noted in the comment column.

In addition, these processors (CPUID signature of DisplayFamily\_DisplayModel 06\_1AH, 06\_1EH, 06\_1FH) also support the following model-specific, product-specific uncore performance monitoring events listed in Table 19-20.

Fixed counters in the core PMU support the architecture events defined in Table 19-2.

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series**

| Event Num. | Umask Value | Event Mask Mnemonic                   | Description   | Comment |
|------------|-------------|---------------------------------------|---|---------|
| 04H        | 07H         | SB_DRAIN.ANY                          | Counts the number of store buffer drains.   |         |
| 06H        | 04H         | STORE_BLOCKS.AT_RET                   | Counts number of loads delayed with at-Retirement block code. The following loads need to be executed at retirement and wait for all senior stores on the same thread to be drained: load splitting across 4K boundary (page split), load accessing uncacheable (UC or WC) memory, load lock, and load with page table in UC or WC memory region. |         |
| 06H        | 08H         | STORE_BLOCKS.L1D_BLOCK                | Cacheable loads delayed with L1D block code.  |         |
| 07H        | 01H         | PARTIAL_ADDRESS_ALIAS                 | Counts false dependency due to partial address aliasing.  |         |
| 08H        | 01H         | DTLB_LOAD_MISSES.ANY                  | Counts all load misses that cause a page walk.  |         |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED       | Counts number of completed page walks due to load miss in the STLB.   |         |
| 08H        | 10H         | DTLB_LOAD_MISSES.STLB_HIT             | Number of cache load STLB hits.   |         |
| 08H        | 20H         | DTLB_LOAD_MISSES.PDE_MISSES           | Number of DTLB cache load misses where the low part of the linear to physical address translation was missed.   |         |
| 08H        | 80H         | DTLB_LOAD_MISSES.LARGE_WALK_COMPLETED | Counts number of completed large page walks due to load miss in the STLB.   |         |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                            | Description  | Comment                                    |
|------------|-------------|--|--|--|
| 0BH        | 01H         | MEM_INST_RETIRED.LOADS                         | Counts the number of instructions with an architecturally-visible load retired on the architected path.  |  |
| 0BH        | 02H         | MEM_INST_RETIRED.STORES                        | Counts the number of instructions with an architecturally-visible store retired on the architected path.   |  |
| 0BH        | 10H         | MEM_INST_RETIRED.LATENCY_ABOVE_THRESHOLD       | Counts the number of instructions exceeding the latency specified with Id_lat facility.  | In conjunction with Id_lat facility.       |
| 0CH        | 01H         | MEM_STORE_RETIRED.DTLB_MISS                    | The event counts the number of retired stores that missed the DTLB. The DTLB miss is not counted if the store operation causes a fault. Does not counter prefetches. Counts both primary and secondary misses to the TLB.                          |  |
| 0EH        | 01H         | UOPS_ISSUED.ANY                                | Counts the number of Uops issued by the Register Allocation Table to the Reservation Station, i.e. the UOPs issued from the front end to the back end.   |  |
| 0EH        | 01H         | UOPS_ISSUED.STALLED_CYCLE_S                    | Counts the number of cycles no Uops issued by the Register Allocation Table to the Reservation Station, i.e. the UOPs issued from the front end to the back end.   | Set "invert=1, cmask = 1".                 |
| 0EH        | 02H         | UOPS_ISSUED.FUSED                              | Counts the number of fused Uops that were issued from the Register Allocation Table to the Reservation Station.  |  |
| 0FH        | 01H         | MEM_UNCORE_RETIRED.L3_DATA_MISS_UNKNOWN        | Counts number of memory load instructions retired where the memory reference missed L3 and data source is unknown.   | Available only for CPUID signature 06_2EH. |
| 0FH        | 02H         | MEM_UNCORE_RETIRED.OTHER_CORE_L2_HITM          | Counts number of memory load instructions retired where the memory reference hit modified data in a sibling core residing on the same socket.  |  |
| 0FH        | 08H         | MEM_UNCORE_RETIRED.REMOTE_CACHE_LOCAL_HOME_HIT | Counts number of memory load instructions retired where the memory reference missed the L1, L2 and L3 caches and HIT in a remote socket's cache. Only counts locally homed lines.  |  |
| 0FH        | 10H         | MEM_UNCORE_RETIRED.REMOTE_DRAM                 | Counts number of memory load instructions retired where the memory reference missed the L1, L2 and L3 caches and was remotely homed. This includes both DRAM access and HITM in a remote socket's cache for remotely homed lines.                  |  |
| 0FH        | 20H         | MEM_UNCORE_RETIRED.LOCAL_DRAM                  | Counts number of memory load instructions retired where the memory reference missed the L1, L2 and L3 caches and required a local socket memory reference. This includes locally homed cachelines that were in a modified state in another socket. |  |
| 0FH        | 80H         | MEM_UNCORE_RETIRED.UNCACHEABLE                 | Counts number of memory load instructions retired where the memory reference missed the L1, L2 and L3 caches and to perform I/O.   | Available only for CPUID signature 06_2EH. |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                  | Description   | Comment |
|------------|-------------|--------------------------------------|---|---------|
| 10H        | 01H         | FP_COMP_OPS_EXE.X87                  | Counts the number of FP Computational Uops Executed. The number of FADD, FSUB, FCOM, FMULs, integer MULs and IMULs, FDIVs, FPREMs, FSQRTS, integer DIVs, and IDIVs. This event does not distinguish an FADD used in the middle of a transcendental flow from a separate FADD instruction. |         |
| 10H        | 02H         | FP_COMP_OPS_EXE.MMX                  | Counts number of MMX Uops executed.   |         |
| 10H        | 04H         | FP_COMP_OPS_EXE.SSE_FP               | Counts number of SSE and SSE2 FP uops executed.   |         |
| 10H        | 08H         | FP_COMP_OPS_EXE.SSE2_INTEGER         | Counts number of SSE2 integer uops executed.  |         |
| 10H        | 10H         | FP_COMP_OPS_EXE.SSE_FP_PACKED        | Counts number of SSE FP packed uops executed.   |         |
| 10H        | 20H         | FP_COMP_OPS_EXE.SSE_FP_SCALAR        | Counts number of SSE FP scalar uops executed.   |         |
| 10H        | 40H         | FP_COMP_OPS_EXE.SSE_SINGLE_PRECISION | Counts number of SSE* FP single precision uops executed.  |         |
| 10H        | 80H         | FP_COMP_OPS_EXE.SSE_DOUBLE_PRECISION | Counts number of SSE* FP double precision uops executed.  |         |
| 12H        | 01H         | SIMD_INT_128.PACKED_MPY              | Counts number of 128 bit SIMD integer multiply operations.  |         |
| 12H        | 02H         | SIMD_INT_128.PACKED_SHIFT            | Counts number of 128 bit SIMD integer shift operations.   |         |
| 12H        | 04H         | SIMD_INT_128.PACK                    | Counts number of 128 bit SIMD integer pack operations.  |         |
| 12H        | 08H         | SIMD_INT_128.UNPACK                  | Counts number of 128 bit SIMD integer unpack operations.  |         |
| 12H        | 10H         | SIMD_INT_128.PACKED_LOGICAL          | Counts number of 128 bit SIMD integer logical operations.   |         |
| 12H        | 20H         | SIMD_INT_128.PACKED_ARITH            | Counts number of 128 bit SIMD integer arithmetic operations.  |         |
| 12H        | 40H         | SIMD_INT_128.SHUFFLE_MOVE            | Counts number of 128 bit SIMD integer shuffle and move operations.  |         |
| 13H        | 01H         | LOAD_DISPATCH.RS                     | Counts number of loads dispatched from the Reservation Station that bypass the Memory Order Buffer.   |         |
| 13H        | 02H         | LOAD_DISPATCH.RS_DELAYED             | Counts the number of delayed RS dispatches at the stage latch. If an RS dispatch cannot bypass to LB, it has another chance to dispatch from the one-cycle delayed staging latch before it is written into the LB.  |         |
| 13H        | 04H         | LOAD_DISPATCH.MOB                    | Counts the number of loads dispatched from the Reservation Station to the Memory Order Buffer.  |         |
| 13H        | 07H         | LOAD_DISPATCH.ANY                    | Counts all loads dispatched from the Reservation Station.   |         |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic     | Description  | Comment  |
|------------|-------------|-------------------------|--|--|
| 14H        | 01H         | ARITH.CYCLES_DIV_BUSY   | Counts the number of cycles the divider is busy executing divide or square root operations. The divide can be integer, X87 or Streaming SIMD Extensions (SSE). The square root operation can be either X87 or SSE.<br>Set 'edge =1, invert=1, cmask=1' to count the number of divides.   | Count may be incorrect When SMT is on.   |
| 14H        | 02H         | ARITH.MUL               | Counts the number of multiply operations executed. This includes integer as well as floating point multiply operations but excludes DPPS mul and MPSAD.  | Count may be incorrect When SMT is on.   |
| 17H        | 01H         | INST_QUEUE_WRITES       | Counts the number of instructions written into the instruction queue every cycle.  |  |
| 18H        | 01H         | INST_DECODED.DECO       | Counts number of instructions that require decoder 0 to be decoded. Usually, this means that the instruction maps to more than 1 uop.  |  |
| 19H        | 01H         | TWO_UOP_INSTS_DECODED   | An instruction that generates two uops was decoded.  |  |
| 1EH        | 01H         | INST_QUEUE_WRITE_CYCLES | This event counts the number of cycles during which instructions are written to the instruction queue. Dividing this counter by the number of instructions written to the instruction queue (INST_QUEUE_WRITES) yields the average number of instructions decoded each cycle. If this number is less than four and the pipe stalls, this indicates that the decoder is failing to decode enough instructions per cycle to sustain the 4-wide pipeline. | If SSE* instructions that are 6 bytes or longer arrive one after another, then front end throughput may limit execution speed. |
| 20H        | 01H         | LSD_OVERFLOW            | Counts number of loops that can't stream from the instruction queue.   |  |
| 24H        | 01H         | L2_RQSTS.LD_HIT         | Counts number of loads that hit the L2 cache. L2 loads include both L1D demand misses as well as L1D prefetches. L2 loads can be rejected for various reasons. Only non rejected loads are counted.  |  |
| 24H        | 02H         | L2_RQSTS.LD_MISS        | Counts the number of loads that miss the L2 cache. L2 loads include both L1D demand misses as well as L1D prefetches.  |  |
| 24H        | 03H         | L2_RQSTS.LOADS          | Counts all L2 load requests. L2 loads include both L1D demand misses as well as L1D prefetches.  |  |
| 24H        | 04H         | L2_RQSTS.RFO_HIT        | Counts the number of store RFO requests that hit the L2 cache. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches. Count includes WC memory requests, where the data is not fetched but the permission to write the line is required.  |  |
| 24H        | 08H         | L2_RQSTS.RFO_MISS       | Counts the number of store RFO requests that miss the L2 cache. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches.  |  |



**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment |
|------------|-------------|--------------------------------|---|---------|
| 24H        | 0CH         | L2_RQSTS.RFOS                  | Counts all L2 store RFO requests. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches.   |         |
| 24H        | 10H         | L2_RQSTS.IFETCH_HIT            | Counts number of instruction fetches that hit the L2 cache. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.  |         |
| 24H        | 20H         | L2_RQSTS.IFETCH_MISS           | Counts number of instruction fetches that miss the L2 cache. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.   |         |
| 24H        | 30H         | L2_RQSTS.IFETCHES              | Counts all instruction fetches. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.  |         |
| 24H        | 40H         | L2_RQSTS.PREFETCH_HIT          | Counts L2 prefetch hits for both code and data.   |         |
| 24H        | 80H         | L2_RQSTS.PREFETCH_MISS         | Counts L2 prefetch misses for both code and data.   |         |
| 24H        | C0H         | L2_RQSTS.PREFETCHES            | Counts all L2 prefetches for both code and data.  |         |
| 24H        | AAH         | L2_RQSTS.MISS                  | Counts all L2 misses for both code and data.  |         |
| 24H        | FFH         | L2_RQSTS.REFERENCES            | Counts all L2 requests for both code and data.  |         |
| 26H        | 01H         | L2_DATA_RQSTS.DEMAND.I_STATE   | Counts number of L2 data demand loads where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss. L2 demand loads are both L1D demand misses and L1D prefetches.               |         |
| 26H        | 02H         | L2_DATA_RQSTS.DEMAND.S_STATE   | Counts number of L2 data demand loads where the cache line to be loaded is in the S (shared) state. L2 demand loads are both L1D demand misses and L1D prefetches.                                    |         |
| 26H        | 04H         | L2_DATA_RQSTS.DEMAND.E_STATE   | Counts number of L2 data demand loads where the cache line to be loaded is in the E (exclusive) state. L2 demand loads are both L1D demand misses and L1D prefetches.                                 |         |
| 26H        | 08H         | L2_DATA_RQSTS.DEMAND.M_STATE   | Counts number of L2 data demand loads where the cache line to be loaded is in the M (modified) state. L2 demand loads are both L1D demand misses and L1D prefetches.                                  |         |
| 26H        | 0FH         | L2_DATA_RQSTS.DEMAND.MESI      | Counts all L2 data demand requests. L2 demand loads are both L1D demand misses and L1D prefetches.  |         |
| 26H        | 10H         | L2_DATA_RQSTS.PREFETCH.I_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss.  |         |
| 26H        | 20H         | L2_DATA_RQSTS.PREFETCH.S_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the S (shared) state. A prefetch RFO will miss on an S state line, while a prefetch read will hit on an S state line. |         |
| 26H        | 40H         | L2_DATA_RQSTS.PREFETCH.E_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the E (exclusive) state.  |         |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment                       |
|------------|-------------|--------------------------------|--|-------------------------------|
| 26H        | 80H         | L2_DATA_RQSTS.PREFETCH.M_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the M (modified) state.  |                               |
| 26H        | F0H         | L2_DATA_RQSTS.PREFETCH.MESI    | Counts all L2 prefetch requests.   |                               |
| 26H        | FFH         | L2_DATA_RQSTS.ANY              | Counts all L2 data requests.   |                               |
| 27H        | 01H         | L2_WRITE.RFO.I_STATE           | Counts number of L2 demand store RFO requests where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss. The L1D prefetcher does not issue a RFO prefetch. | This is a demand RFO request. |
| 27H        | 02H         | L2_WRITE.RFO.S_STATE           | Counts number of L2 store RFO requests where the cache line to be loaded is in the S (shared) state. The L1D prefetcher does not issue a RFO prefetch.                             | This is a demand RFO request. |
| 27H        | 08H         | L2_WRITE.RFO.M_STATE           | Counts number of L2 store RFO requests where the cache line to be loaded is in the M (modified) state. The L1D prefetcher does not issue a RFO prefetch.                           | This is a demand RFO request. |
| 27H        | 0EH         | L2_WRITE.RFO.HIT               | Counts number of L2 store RFO requests where the cache line to be loaded is in either the S, E or M states. The L1D prefetcher does not issue a RFO prefetch.                      | This is a demand RFO request. |
| 27H        | 0FH         | L2_WRITE.RFO.MESI              | Counts all L2 store RFO requests. The L1D prefetcher does not issue a RFO prefetch.  | This is a demand RFO request. |
| 27H        | 10H         | L2_WRITE.LOCK.I_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the I (invalid) state, for example, a cache miss.   |                               |
| 27H        | 20H         | L2_WRITE.LOCK.S_STATE          | Counts number of L2 lock RFO requests where the cache line to be loaded is in the S (shared) state.  |                               |
| 27H        | 40H         | L2_WRITE.LOCK.E_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the E (exclusive) state.  |                               |
| 27H        | 80H         | L2_WRITE.LOCK.M_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the M (modified) state.   |                               |
| 27H        | E0H         | L2_WRITE.LOCK.HIT              | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in either the S, E, or M state.  |                               |
| 27H        | F0H         | L2_WRITE.LOCK.MESI             | Counts all L2 demand lock RFO requests.  |                               |
| 28H        | 01H         | L1D_WB_L2.I_STATE              | Counts number of L1 writebacks to the L2 where the cache line to be written is in the I (invalid) state, i.e., a cache miss.   |                               |
| 28H        | 02H         | L1D_WB_L2.S_STATE              | Counts number of L1 writebacks to the L2 where the cache line to be written is in the S state.   |                               |
| 28H        | 04H         | L1D_WB_L2.E_STATE              | Counts number of L1 writebacks to the L2 where the cache line to be written is in the E (exclusive) state.   |                               |
| 28H        | 08H         | L1D_WB_L2.M_STATE              | Counts number of L1 writebacks to the L2 where the cache line to be written is in the M (modified) state.  |                               |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description  | Comment  |
|------------|-------------|---------------------------|--|--|
| 28H        | 0FH         | L1D_WB_L2.MESI            | Counts all L1 writebacks to the L2 .   |  |
| 2EH        | 4FH         | L3_LAT_CACHE.REFERENCE    | This event counts requests originating from the core that reference a cache line in the last level cache. The event count includes speculative traffic but excludes cache line fills due to a L2 hardware-prefetch. Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended. | See Table 19-1.  |
| 2EH        | 41H         | L3_LAT_CACHE.MISS         | This event counts each cache miss condition for references to the last level cache. The event count may include speculative traffic but excludes cache line fills due to L2 hardware-prefetches. Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended.                    | See Table 19-1.  |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling.   | See Table 19-1.  |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF_P    | Increments at the frequency of TSC when not halted.  | See Table 19-1.  |
| 40H        | 01H         | L1D_CACHE_LD.I_STATE      | Counts L1 data cache read requests where the cache line to be loaded is in the I (invalid) state, i.e. the read request missed the cache.  | Counter 0, 1 only.   |
| 40H        | 02H         | L1D_CACHE_LD.S_STATE      | Counts L1 data cache read requests where the cache line to be loaded is in the S (shared) state.   | Counter 0, 1 only.   |
| 40H        | 04H         | L1D_CACHE_LD.E_STATE      | Counts L1 data cache read requests where the cache line to be loaded is in the E (exclusive) state.  | Counter 0, 1 only.   |
| 40H        | 08H         | L1D_CACHE_LD.M_STATE      | Counts L1 data cache read requests where the cache line to be loaded is in the M (modified) state.   | Counter 0, 1 only.   |
| 40H        | 0FH         | L1D_CACHE_LD.MESI         | Counts L1 data cache read requests.  | Counter 0, 1 only.   |
| 41H        | 02H         | L1D_CACHE_ST.S_STATE      | Counts L1 data cache store RFO requests where the cache line to be loaded is in the S (shared) state.  | Counter 0, 1 only.   |
| 41H        | 04H         | L1D_CACHE_ST.E_STATE      | Counts L1 data cache store RFO requests where the cache line to be loaded is in the E (exclusive) state.   | Counter 0, 1 only.   |
| 41H        | 08H         | L1D_CACHE_ST.M_STATE      | Counts L1 data cache store RFO requests where cache line to be loaded is in the M (modified) state.  | Counter 0, 1 only.   |
| 42H        | 01H         | L1D_CACHE_LOCK.HIT        | Counts retired load locks that hit in the L1 data cache or hit in an already allocated fill buffer. The lock portion of the load lock transaction must hit in the L1D.   | The initial load will pull the lock into the L1 data cache. Counter 0, 1 only. |
| 42H        | 02H         | L1D_CACHE_LOCK.S_STATE    | Counts L1 data cache retired load locks that hit the target cache line in the shared state.  | Counter 0, 1 only.   |
| 42H        | 04H         | L1D_CACHE_LOCK.E_STATE    | Counts L1 data cache retired load locks that hit the target cache line in the exclusive state.   | Counter 0, 1 only.   |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description   | Comment  |
|------------|-------------|----------------------------------|---|--|
| 42H        | 08H         | L1D_CACHE_LOCK.M_STATE           | Counts L1 data cache retired load locks that hit the target cache line in the modified state.   | Counter 0, 1 only.   |
| 43H        | 01H         | L1D_ALL_REF.ANY                  | Counts all references (uncached, speculated and retired) to the L1 data cache, including all loads and stores with any memory types. The event counts memory accesses only when they are actually performed. For example, a load blocked by unknown store address and later performed is only counted once.   | The event does not include non-memory accesses, such as I/O accesses. Counter 0, 1 only. |
| 43H        | 02H         | L1D_ALL_REF.CACHEABLE            | Counts all data reads and writes (speculated and retired) from cacheable memory, including locked operations.   | Counter 0, 1 only.   |
| 49H        | 01H         | DTLB_MISSES.ANY                  | Counts the number of misses in the STLB which causes a page walk.   |  |
| 49H        | 02H         | DTLB_MISSES.WALK_COMPLETED       | Counts number of misses in the STLB which resulted in a completed page walk.  |  |
| 49H        | 10H         | DTLB_MISSES.STLB_HIT             | Counts the number of DTLB first level misses that hit in the second level TLB. This event is only relevant if the core contains multiple DTLB levels.   |  |
| 49H        | 20H         | DTLB_MISSES.PDE_MISS             | Number of DTLB misses caused by low part of address, includes references to 2M pages because 2M pages do not use the PDE.   |  |
| 49H        | 80H         | DTLB_MISSES.LARGE_WALK_COMPLETED | Counts number of misses in the STLB which resulted in a completed page walk for large pages.  |  |
| 4CH        | 01H         | LOAD_HIT_PRE                     | Counts load operations sent to the L1 data cache while a previous SSE prefetch instruction to the same cache line has started prefetching but has not yet finished.   |  |
| 4EH        | 01H         | L1D_PREFETCH.REQUESTS            | Counts number of hardware prefetch requests dispatched out of the prefetch FIFO.  |  |
| 4EH        | 02H         | L1D_PREFETCH.MISS                | Counts number of hardware prefetch requests that miss the L1D. There are two prefetchers in the L1D. A streamer, which predicts lines sequentially after this one should be fetched, and the IP prefetcher that remembers access patterns for the current instruction. The streamer prefetcher stops on an L1D hit, while the IP prefetcher does not. |  |
| 4EH        | 04H         | L1D_PREFETCH.TRIGGERS            | Counts number of prefetch requests triggered by the Finite State Machine and pushed into the prefetch FIFO. Some of the prefetch requests are dropped due to overwrites or competition between the IP index prefetcher and streamer prefetcher. The prefetch FIFO contains 4 entries.   |  |
| 51H        | 01H         | L1D.REPL                         | Counts the number of lines brought into the L1 data cache.  | Counter 0, 1 only.   |
| 51H        | 02H         | L1D.M_REPL                       | Counts the number of modified lines brought into the L1 data cache.   | Counter 0, 1 only.   |
| 51H        | 04H         | L1D.M_EVICT                      | Counts the number of modified lines evicted from the L1 data cache due to replacement.  | Counter 0, 1 only.   |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment   |
|------------|-------------|--------------------------------|--|---|
| 51H        | 08H         | L1D.M_SNOOP_EVICT              | Counts the number of modified lines evicted from the L1 data cache due to snoop HITM intervention.   | Counter 0, 1 only.  |
| 52H        | 01H         | L1D_CACHE_PREFETCH_LOCK_FB_HIT | Counts the number of cacheable load lock speculated instructions accepted into the fill buffer.  |   |
| 53H        | 01H         | L1D_CACHE_LOCK_FB_HIT          | Counts the number of cacheable load lock speculated or retired instructions accepted into the fill buffer.   |   |
| 63H        | 01H         | CACHE_LOCK_CYCLES.L1D_L2       | Cycle count during which the L1D and L2 are locked. A lock is asserted when there is a locked memory access, due to uncacheable memory, a locked operation that spans two cache lines, or a page walk from an uncacheable page table.                                | Counter 0, 1 only. L1D and L2 locks have a very high performance penalty and it is highly recommended to avoid such accesses. |
| 63H        | 02H         | CACHE_LOCK_CYCLES.L1D          | Counts the number of cycles that cacheline in the L1 data cache unit is locked.  | Counter 0, 1 only.  |
| 6CH        | 01H         | IO_TRANSACTIONS                | Counts the number of completed I/O transactions.   |   |
| 80H        | 01H         | L1I.HITS                       | Counts all instruction fetches that hit the L1 instruction cache.  |   |
| 80H        | 02H         | L1I.MISSES                     | Counts all instruction fetches that miss the L1I cache. This includes instruction cache misses, streaming buffer misses, victim cache misses and uncacheable fetches. An instruction fetch miss is counted only once and not once for every cycle it is outstanding. |   |
| 80H        | 03H         | L1I.READS                      | Counts all instruction fetches, including uncacheable fetches that bypass the L1I.   |   |
| 80H        | 04H         | L1I.CYCLES_STALLED             | Cycle counts for which an instruction fetch stalls due to a L1I cache miss, ITLB miss or ITLB fault.   |   |
| 82H        | 01H         | LARGE_ITLB.HIT                 | Counts number of large ITLB hits.  |   |
| 85H        | 01H         | ITLB_MISSES.ANY                | Counts the number of misses in all levels of the ITLB which causes a page walk.  |   |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETED     | Counts number of misses in all levels of the ITLB which resulted in a completed page walk.   |   |
| 87H        | 01H         | ILD_STALL.LCP                  | Cycles Instruction Length Decoder stalls due to length changing prefixes: 66, 67 or REX.W (for Intel 64) instructions which change the length of the decoded instruction.  |   |
| 87H        | 02H         | ILD_STALL.MRU                  | Instruction Length Decoder stall cycles due to Branch Prediction Unit (PBU) Most Recently Used (MRU) bypass.   |   |
| 87H        | 04H         | ILD_STALL.IQ_FULL              | Stall cycles due to a full instruction queue.  |   |
| 87H        | 08H         | ILD_STALL.REGEN                | Counts the number of regen stalls.   |   |
| 87H        | 0FH         | ILD_STALL.ANY                  | Counts any cycles the Instruction Length Decoder is stalled.   |   |
| 88H        | 01H         | BR_INST_EXEC.COND              | Counts the number of conditional near branch instructions executed, but not necessarily retired.   |   |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description   | Comment |
|------------|-------------|---------------------------------|---|---------|
| 88H        | 02H         | BR_INST_EXEC.DIRECT             | Counts all unconditional near branch instructions excluding calls and indirect branches.  |         |
| 88H        | 04H         | BR_INST_EXEC.INDIRECT_NON_CALL  | Counts the number of executed indirect near branch instructions that are not calls.   |         |
| 88H        | 07H         | BR_INST_EXEC.NON_CALLS          | Counts all non-call near branch instructions executed, but not necessarily retired.   |         |
| 88H        | 08H         | BR_INST_EXEC.RETURN_NEAR        | Counts indirect near branches that have a return mnemonic.  |         |
| 88H        | 10H         | BR_INST_EXEC.DIRECT_NEAR_CALL   | Counts unconditional near call branch instructions, excluding non-call branch, executed.  |         |
| 88H        | 20H         | BR_INST_EXEC.INDIRECT_NEAR_CALL | Counts indirect near calls, including both register and memory indirect, executed.  |         |
| 88H        | 30H         | BR_INST_EXEC.NEAR_CALLS         | Counts all near call branches executed, but not necessarily retired.  |         |
| 88H        | 40H         | BR_INST_EXEC.TAKEN              | Counts taken near branches executed, but not necessarily retired.   |         |
| 88H        | 7FH         | BR_INST_EXEC.ANY                | Counts all near executed branches (not necessarily retired). This includes only instructions and not micro-op branches. Frequent branching is not necessarily a major performance issue. However frequent branch mispredictions may be a problem. |         |
| 89H        | 01H         | BR_MISP_EXEC.COND               | Counts the number of mispredicted conditional near branch instructions executed, but not necessarily retired.   |         |
| 89H        | 02H         | BR_MISP_EXEC.DIRECT             | Counts mispredicted macro unconditional near branch instructions, excluding calls and indirect branches (should always be 0).   |         |
| 89H        | 04H         | BR_MISP_EXEC.INDIRECT_NON_CALL  | Counts the number of executed mispredicted indirect near branch instructions that are not calls.  |         |
| 89H        | 07H         | BR_MISP_EXEC.NON_CALLS          | Counts mispredicted non-call near branches executed, but not necessarily retired.   |         |
| 89H        | 08H         | BR_MISP_EXEC.RETURN_NEAR        | Counts mispredicted indirect branches that have a return mnemonic.  |         |
| 89H        | 10H         | BR_MISP_EXEC.DIRECT_NEAR_CALL   | Counts mispredicted non-indirect near calls executed, (should always be 0).   |         |
| 89H        | 20H         | BR_MISP_EXEC.INDIRECT_NEAR_CALL | Counts mispredicted indirect near calls executed, including both register and memory indirect.  |         |
| 89H        | 30H         | BR_MISP_EXEC.NEAR_CALLS         | Counts all mispredicted near call branches executed, but not necessarily retired.   |         |
| 89H        | 40H         | BR_MISP_EXEC.TAKEN              | Counts executed mispredicted near branches that are taken, but not necessarily retired.   |         |
| 89H        | 7FH         | BR_MISP_EXEC.ANY                | Counts the number of mispredicted near branch instructions that were executed, but not necessarily retired.   |         |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic         | Description  | Comment   |
|------------|-------------|-----------------------------|--|---|
| A2H        | 01H         | RESOURCE_STALLS.ANY         | Counts the number of Allocator resource related stalls. Includes register renaming buffer entries, memory buffer entries. In addition to resource related stalls, this event counts some other events. Includes stalls arising during branch misprediction recovery, such as if retirement of the mispredicted branch is delayed and stalls arising while store buffer is draining from synchronizing operations.  | Does not include stalls due to SuperQ (off core) queue full, too many cache misses, etc.    |
| A2H        | 02H         | RESOURCE_STALLS.LOAD        | Counts the cycles of stall due to lack of load buffer for load operation.  |   |
| A2H        | 04H         | RESOURCE_STALLS.RS_FULL     | This event counts the number of cycles when the number of instructions in the pipeline waiting for execution reaches the limit the processor can handle. A high count of this event indicates that there are long latency operations in the pipe (possibly load and store operations that miss the L2 cache, or instructions dependent upon instructions further down the pipeline that have yet to retire.  | When RS is full, new instructions cannot enter the reservation station and start execution. |
| A2H        | 08H         | RESOURCE_STALLS.STORE       | This event counts the number of cycles that a resource related stall will occur due to the number of store instructions reaching the limit of the pipeline, (i.e. all store buffers are used). The stall ends when a store instruction commits its data to the cache or memory.  |   |
| A2H        | 10H         | RESOURCE_STALLS.ROB_FULL    | Counts the cycles of stall due to re-order buffer full.  |   |
| A2H        | 20H         | RESOURCE_STALLS.FPCW        | Counts the number of cycles while execution was stalled due to writing the floating-point unit (FPU) control word.   |   |
| A2H        | 40H         | RESOURCE_STALLS.MXCSR       | Stalls due to the MXCSR register rename occurring to close to a previous MXCSR rename. The MXCSR provides control and status for the MMX registers.  |   |
| A2H        | 80H         | RESOURCE_STALLS.OTHER       | Counts the number of cycles while execution was stalled due to other resource issues.  |   |
| A6H        | 01H         | MACRO_INSTS.FUSIONS_DECODED | Counts the number of instructions decoded that are macro-fused but not necessarily executed or retired.  |   |
| A7H        | 01H         | BACLEAR_FORCE_IQ            | Counts number of times a BACLEAR was forced by the Instruction Queue. The IQ is also responsible for providing conditional branch prediction direction based on a static scheme and dynamic data provided by the L2 Branch Prediction Unit. If the conditional branch target is not found in the Target Array and the IQ predicts that the branch is taken, then the IQ will force the Branch Address Calculator to issue a BACLEAR. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline. |   |
| A8H        | 01H         | LSD.UOPS                    | Counts the number of micro-ops delivered by loop stream detector.  | Use cmask=1 and invert to count cycles.   |
| AEH        | 01H         | ITLB_FLUSH                  | Counts the number of ITLB flushes.   |   |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description  | Comment                                      |
|------------|-------------|---|--|--|
| B0H        | 40H         | OFFCORE_REQUESTS.L1D_WRITEBACK            | Counts number of L1D writebacks to the uncore.   |  |
| B1H        | 01H         | UOPS_EXECUTED.PORT0                       | Counts number of uops executed that were issued on port 0. Port 0 handles integer arithmetic, SIMD and FP add uops.  |  |
| B1H        | 02H         | UOPS_EXECUTED.PORT1                       | Counts number of uops executed that were issued on port 1. Port 1 handles integer arithmetic, SIMD, integer shift, FP multiply and FP divide uops.   |  |
| B1H        | 04H         | UOPS_EXECUTED.PORT2_CORE                  | Counts number of uops executed that were issued on port 2. Port 2 handles the load uops. This is a core count only and cannot be collected per thread.   |  |
| B1H        | 08H         | UOPS_EXECUTED.PORT3_CORE                  | Counts number of uops executed that were issued on port 3. Port 3 handles store uops. This is a core count only and cannot be collected per thread.  |  |
| B1H        | 10H         | UOPS_EXECUTED.PORT4_CORE                  | Counts number of uops executed that where issued on port 4. Port 4 handles the value to be stored for the store uops issued on port 3. This is a core count only and cannot be collected per thread.   |  |
| B1H        | 1FH         | UOPS_EXECUTED.CORE_ACTIVE_CYCLES_NO_PORT5 | Counts cycles when the uops executed were issued from any ports except port 5. Use Cmask=1 for active cycles; Cmask=0 for weighted cycles. Use CMask=1, Invert=1 to count P0-4 stalled cycles. Use Cmask=1, Edge=1, Invert=1 to count P0-4 stalls. |  |
| B1H        | 20H         | UOPS_EXECUTED.PORT5                       | Counts number of uops executed that where issued on port 5.  |  |
| B1H        | 3FH         | UOPS_EXECUTED.CORE_ACTIVE_CYCLES          | Counts cycles when the uops are executing. Use Cmask=1 for active cycles; Cmask=0 for weighted cycles. Use CMask=1, Invert=1 to count P0-4 stalled cycles. Use Cmask=1, Edge=1, Invert=1 to count P0-4 stalls.                                     |  |
| B1H        | 40H         | UOPS_EXECUTED.PORT015                     | Counts number of uops executed that where issued on port 0, 1, or 5.   | Use cmask=1, invert=1 to count stall cycles. |
| B1H        | 80H         | UOPS_EXECUTED.PORT234                     | Counts number of uops executed that where issued on port 2, 3, or 4.   |  |
| B2H        | 01H         | OFFCORE_REQUESTS_SQ_FULL                  | Counts number of cycles the SQ is full to handle off-core requests.  |  |
| B7H        | 01H         | OFF_CORE_RESPONSE_0                       | See Section 18.3.1.1.3, "Off-core Response Performance Monitoring in the Processor Core".  | Requires programming MSR 01A6H.              |
| B8H        | 01H         | SNOOP_RESPONSE.HIT                        | Counts HIT snoop response sent by this thread in response to a snoop request.  |  |
| B8H        | 02H         | SNOOP_RESPONSE.HITE                       | Counts HIT E snoop response sent by this thread in response to a snoop request.  |  |
| B8H        | 04H         | SNOOP_RESPONSE.HITM                       | Counts HIT M snoop response sent by this thread in response to a snoop request.  |  |
| BBH        | 01H         | OFF_CORE_RESPONSE_1                       | See Section 18.6.3, "Performance Monitoring (Processors Based on Intel NetBurst® Microarchitecture)".  | Requires programming MSR 01A7H.              |



**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description   | Comment  |
|------------|-------------|---------------------------------|---|--|
| C0H        | 00H         | INST_RETIRED.ANY_P              | See Table 19-1.<br>Notes: INST_RETIRED.ANY is counted by a designated fixed counter. INST_RETIRED.ANY_P is counted by a programmable counter and is an architectural performance event. Event is supported if CPUID.A.EBX[1] = 0.   | Counting: Faulting executions of GETSEC/VM entry/VM Exit/MWait will not count as retired instructions. |
| C0H        | 02H         | INST_RETIRED.X87                | Counts the number of MMX instructions retired.  |  |
| C0H        | 04H         | INST_RETIRED.MMX                | Counts the number of floating point computational operations retired: floating point computational operations executed by the assist handler and sub-operations of complex floating point instructions like transcendental instructions.  |  |
| C2H        | 01H         | UOPS_RETIRED.ANY                | Counts the number of micro-ops retired, (macro-fused=1, micro-fused=2, others=1; maximum count of 8 per cycle). Most instructions are composed of one or two micro-ops. Some instructions are decoded into longer sequences such as repeat instructions, floating point transcendental instructions, and assists. | Use cmask=1 and invert to count active cycles or stalled cycles.                                       |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOTS       | Counts the number of retirement slots used each cycle.  |  |
| C2H        | 04H         | UOPS_RETIRED.MACRO_FUSED        | Counts number of macro-fused uops retired.  |  |
| C3H        | 01H         | MACHINE_CLEAR.CYCLES            | Counts the cycles machine clear is asserted.  |  |
| C3H        | 02H         | MACHINE_CLEAR.MEM_ORDER         | Counts the number of machine clears due to memory order conflicts.  |  |
| C3H        | 04H         | MACHINE_CLEAR.SMC               | Counts the number of times that a program writes to a code section. Self-modifying code causes a severe penalty in all Intel 64 and IA-32 processors. The modified cache line is written back to the L2 and L3caches.   |  |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES    | Branch instructions at retirement.  | See Table 19-1.  |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL     | Counts the number of conditional branch instructions retired.   |  |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL       | Counts the number of direct & indirect near unconditional calls retired.  |  |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES    | Mispredicted branch instructions at retirement.   | See Table 19-1.  |
| C5H        | 02H         | BR_MISP_RETIRED.NEAR_CALL       | Counts mispredicted direct & indirect near unconditional retired calls.   |  |
| C7H        | 01H         | SSEX_UOPS_RETIRED.PACKED_SINGLE | Counts SIMD packed single-precision floating point Uops retired.  |  |
| C7H        | 02H         | SSEX_UOPS_RETIRED.SCALAR_SINGLE | Counts SIMD scalar single-precision floating point Uops retired.  |  |
| C7H        | 04H         | SSEX_UOPS_RETIRED.PACKED_DOUBLE | Counts SIMD packed double-precision floating point Uops retired.  |  |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment |
|------------|-------------|---|---|---------|
| C7H        | 08H         | SSEX_UOPS_RETIRED.SCALAR_DOUBLE         | Counts SIMD scalar double-precision floating point Uops retired.  |         |
| C7H        | 10H         | SSEX_UOPS_RETIRED.VECTOR_INTEGER        | Counts 128-bit SIMD vector integer Uops retired.  |         |
| C8H        | 20H         | ITLB_MISS_RETIRED                       | Counts the number of retired instructions that missed the ITLB when the instruction was fetched.  |         |
| CBH        | 01H         | MEM_LOAD_RETIRED.L1D_HIT                | Counts number of retired loads that hit the L1 data cache.  |         |
| CBH        | 02H         | MEM_LOAD_RETIRED.L2_HIT                 | Counts number of retired loads that hit the L2 data cache.  |         |
| CBH        | 04H         | MEM_LOAD_RETIRED.L3_UNSHARED_HIT        | Counts number of retired loads that hit their own, unshared lines in the L3 cache.  |         |
| CBH        | 08H         | MEM_LOAD_RETIRED.OTHER_CORE_L2_HIT_HITM | Counts number of retired loads that hit in a sibling core's L2 (on die core). Since the L3 is inclusive of all cores on the package, this is an L3 hit. This counts both clean and modified hits.   |         |
| CBH        | 10H         | MEM_LOAD_RETIRED.L3_MISS                | Counts number of retired loads that miss the L3 cache. The load was satisfied by a remote socket, local memory or an IOH.   |         |
| CBH        | 40H         | MEM_LOAD_RETIRED.HIT_LFB                | Counts number of retired loads that miss the L1D and the address is located in an allocated line fill buffer and will soon be committed to cache. This is counting secondary L1D misses.  |         |
| CBH        | 80H         | MEM_LOAD_RETIRED.DTLB_MISS              | Counts the number of retired loads that missed the DTLB. The DTLB miss is not counted if the load operation causes a fault. This event counts loads from cacheable memory only. The event does not count loads by software prefetches. Counts both primary and secondary misses to the TLB. |         |
| CCH        | 01H         | FP_MMX_TRANS.TO_FP                      | Counts the first floating-point instruction following any MMX instruction. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |         |
| CCH        | 02H         | FP_MMX_TRANS.TO_MMX                     | Counts the first MMX instruction following a floating-point instruction. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |         |
| CCH        | 03H         | FP_MMX_TRANS.ANY                        | Counts all transitions from floating point to MMX instructions and from MMX instructions to floating point instructions. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |         |
| D0H        | 01H         | MACRO_INSTS.DECODED                     | Counts the number of instructions decoded, (but not necessarily executed or retired).   |         |
| D1H        | 02H         | UOPS_DECODED.MS                         | Counts the number of Uops decoded by the Microcode Sequencer, MS. The MS delivers uops when the instruction is more than 4 uops long or a microcode assist is occurring.  |         |

**Table 19-19. Performance Events In the Processor Core for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic      | Description   | Comment |
|------------|-------------|--------------------------|---|---------|
| D1H        | 04H         | UOPS_DECODED.ESP_FOLDING | Counts number of stack pointer (ESP) instructions decoded: push, pop, call, ret, etc. ESP instructions do not generate a Uop to increment or decrement ESP. Instead, they update an ESP_Offset register that keeps track of the delta to the current value of the ESP register.   |         |
| D1H        | 08H         | UOPS_DECODED.ESP_SYNC    | Counts number of stack pointer (ESP) sync operations where an ESP instruction is corrected by adding the ESP offset register to the current value of the ESP register.  |         |
| D2H        | 01H         | RAT_STALLS.FLAGS         | Counts the number of cycles during which execution stalled due to several reasons, one of which is a partial flag register stall. A partial register stall may occur when two conditions are met: 1) an instruction modifies some, but not all, of the flags in the flag register and 2) the next instruction, which depends on flags, depends on flags that were not modified by this instruction.   |         |
| D2H        | 02H         | RAT_STALLS.REGISTERS     | This event counts the number of cycles instruction execution latency became longer than the defined latency because the instruction used a register that was partially written by previous instruction.   |         |
| D2H        | 04H         | RAT_STALLS.ROB_READ_PORT | Counts the number of cycles when ROB read port stalls occurred, which did not allow new micro-ops to enter the out-of-order pipeline. Note that, at this stage in the pipeline, additional stalls may occur at the same cycle and prevent the stalled micro-ops from entering the pipe. In such a case, micro-ops retry entering the execution pipe in the next cycle and the ROB-read port stall is counted again.   |         |
| D2H        | 08H         | RAT_STALLS.SCOREBOARD    | Counts the cycles where we stall due to microarchitecturally required serialization. Microcode scoreboarding stalls.  |         |
| D2H        | 0FH         | RAT_STALLS.ANY           | Counts all Register Allocation Table stall cycles due to: Cycles when ROB read port stalls occurred, which did not allow new micro-ops to enter the execution pipe. Cycles when partial register stalls occurred. Cycles when flag stalls occurred. Cycles floating-point unit (FPU) status word stalls occurred. To count each of these conditions separately use the events: RAT_STALLS.ROB_READ_PORT, RAT_STALLS.PARTIAL, RAT_STALLS.FLAGS, and RAT_STALLS.FPSW. |         |
| D4H        | 01H         | SEG_RENAME_STALLS        | Counts the number of stall cycles due to the lack of renaming resources for the ES, DS, FS, and GS segment registers. If a segment is renamed but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.  |         |
| D5H        | 01H         | ES_REG_RENAMES           | Counts the number of times the ES segment register is renamed.  |         |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic       | Description  | Comment   |
|------------|-------------|---------------------------|--|---|
| DBH        | 01H         | UOP_UNFUSION              | Counts unfusion events due to floating-point exception to a fused uop.   |   |
| E0H        | 01H         | BR_INST_DECODED           | Counts the number of branch instructions decoded.  |   |
| E5H        | 01H         | BPU_MISSED_CALL_RET       | Counts number of times the Branch Prediction Unit missed predicting a call or return branch.   |   |
| E6H        | 01H         | BACLEAR.CLEAR             | Counts the number of times the front end is resteeered, mainly when the Branch Prediction Unit cannot provide a correct prediction and this is corrected by the Branch Address Calculator at the front end. This can occur if the code has many branches such that they cannot be consumed by the BPU. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline. The effect on total execution time depends on the surrounding code. |   |
| E6H        | 02H         | BACLEAR.BAD_TARGET        | Counts number of Branch Address Calculator clears (BACLEAR) asserted due to conditional branch instructions in which there was a target hit but the direction was wrong. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline.   |   |
| E8H        | 01H         | BPU_CLEARS.EARLY          | Counts early (normal) Branch Prediction Unit clears: BPU predicted a taken branch after incorrectly assuming that it was not taken.  | The BPU clear leads to 2 cycle bubble in the front end. |
| E8H        | 02H         | BPU_CLEARS.LATE           | Counts late Branch Prediction Unit clears due to Most Recently Used conflicts. The PBU clear leads to a 3 cycle bubble in the front end.   |   |
| F0H        | 01H         | L2_TRANSACTIONS.LOAD      | Counts L2 load operations due to HW prefetch or demand loads.  |   |
| F0H        | 02H         | L2_TRANSACTIONS.RFO       | Counts L2 RFO operations due to HW prefetch or demand RFOs.  |   |
| F0H        | 04H         | L2_TRANSACTIONS.IFETCH    | Counts L2 instruction fetch operations due to HW prefetch or demand ifetch.  |   |
| F0H        | 08H         | L2_TRANSACTIONS.PREFETCH  | Counts L2 prefetch operations.   |   |
| F0H        | 10H         | L2_TRANSACTIONS.L1D_WB    | Counts L1D writeback operations to the L2.   |   |
| F0H        | 20H         | L2_TRANSACTIONS.FILL      | Counts L2 cache line fill operations due to load, RFO, L1D writeback or prefetch.  |   |
| F0H        | 40H         | L2_TRANSACTIONS.WB        | Counts L2 writeback operations to the L3.  |   |
| F0H        | 80H         | L2_TRANSACTIONS.ANY       | Counts all L2 cache operations.  |   |
| F1H        | 02H         | L2_LINES_IN.S_STATE       | Counts the number of cache lines allocated in the L2 cache in the S (shared) state.  |   |
| F1H        | 04H         | L2_LINES_IN.E_STATE       | Counts the number of cache lines allocated in the L2 cache in the E (exclusive) state.   |   |
| F1H        | 07H         | L2_LINES_IN.ANY           | Counts the number of cache lines allocated in the L2 cache.  |   |
| F2H        | 01H         | L2_LINES_OUT.DEMAND_CLEAN | Counts L2 clean cache lines evicted by a demand request.   |   |

**Table 19-19. Performance Events In the Processor Core for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic         | Description  | Comment |
|------------|-------------|-----------------------------|--|---------|
| F2H        | 02H         | L2_LINES_OUT.DEMAND_DIRTY   | Counts L2 dirty (modified) cache lines evicted by a demand request.  |         |
| F2H        | 04H         | L2_LINES_OUT.PREFETCH_CLEAN | Counts L2 clean cache line evicted by a prefetch request.  |         |
| F2H        | 08H         | L2_LINES_OUT.PREFETCH_DIRTY | Counts L2 modified cache line evicted by a prefetch request.   |         |
| F2H        | 0FH         | L2_LINES_OUT.ANY            | Counts all L2 cache lines evicted for any reason.  |         |
| F4H        | 10H         | SQ_MISC.SPLIT_LOCK          | Counts the number of SQ lock splits across a cache line.   |         |
| F6H        | 01H         | SQ_FULL_STALL_CYCLES        | Counts cycles the Super Queue is full. Neither of the threads on this core will be able to access the uncore.  |         |
| F7H        | 01H         | FP_ASSIST.ALL               | Counts the number of floating point operations executed that required micro-code assist intervention. Assists are required in the following cases: SSE instructions (denormal input when the DAZ flag is off or underflow result when the FTZ flag is off); x87 instructions (NaN or denormal are loaded to a register or used as input from memory, division by 0 or underflow output). |         |
| F7H        | 02H         | FP_ASSIST.OUTPUT            | Counts number of floating point micro-code assist when the output value (destination register) is invalid.   |         |
| F7H        | 04H         | FP_ASSIST.INPUT             | Counts number of floating point micro-code assist when the input value (one of the source operands to an FP instruction) is invalid.   |         |
| FDH        | 01H         | SIMD_INT_64.PACKED_MPY      | Counts number of SIMD integer 64 bit packed multiply operations.   |         |
| FDH        | 02H         | SIMD_INT_64.PACKED_SHIFT    | Counts number of SIMD integer 64 bit packed shift operations.  |         |
| FDH        | 04H         | SIMD_INT_64.PACK            | Counts number of SIMD integer 64 bit pack operations.  |         |
| FDH        | 08H         | SIMD_INT_64.UNPACK          | Counts number of SIMD integer 64 bit unpack operations.  |         |
| FDH        | 10H         | SIMD_INT_64.PACKED_LOGICAL  | Counts number of SIMD integer 64 bit logical operations.   |         |
| FDH        | 20H         | SIMD_INT_64.PACKED_ARITH    | Counts number of SIMD integer 64 bit arithmetic operations.  |         |
| FDH        | 40H         | SIMD_INT_64.SHUFFLE_MOVE    | Counts number of SIMD integer 64 bit shift or move operations.   |         |

Model-specific performance monitoring events that are located in the uncore sub-system are implementation specific between different platforms using processors based on Intel microarchitecture code name Nehalem. Processors with CPUID signature of DisplayFamily\_DisplayModel 06\_1AH, 06\_1EH, and 06\_1FH support performance events listed in Table 19-20.

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series**

| Event Num. | Umask Value | Event Mask Mnemonic                        | Description  | Comment |
|------------|-------------|--|--|---------|
| 00H        | 01H         | UNC_GQ_CYCLES_FULL.READ_TRACKER            | Uncore cycles Global Queue read tracker is full.   |         |
| 00H        | 02H         | UNC_GQ_CYCLES_FULL.WRITE_TRACKER           | Uncore cycles Global Queue write tracker is full.  |         |
| 00H        | 04H         | UNC_GQ_CYCLES_FULL.PEER_PROBE_TRACKER      | Uncore cycles Global Queue peer probe tracker is full. The peer probe tracker queue tracks snoops from the IOH and remote sockets.   |         |
| 01H        | 01H         | UNC_GQ_CYCLES_NOT_EMPTY.READ_TRACKER       | Uncore cycles were Global Queue read tracker has at least one valid entry.   |         |
| 01H        | 02H         | UNC_GQ_CYCLES_NOT_EMPTY.WRITE_TRACKER      | Uncore cycles were Global Queue write tracker has at least one valid entry.  |         |
| 01H        | 04H         | UNC_GQ_CYCLES_NOT_EMPTY.PEER_PROBE_TRACKER | Uncore cycles were Global Queue peer probe tracker has at least one valid entry. The peer probe tracker queue tracks IOH and remote socket snoops.   |         |
| 03H        | 01H         | UNC_GQ_ALLOC.READ_TRACKER                  | Counts the number of read tracker allocate to deallocate entries. The GQ read tracker allocate to deallocate occupancy count is divided by the count to obtain the average read tracker latency.   |         |
| 03H        | 02H         | UNC_GQ_ALLOC.RT_L3_MISS                    | Counts the number GQ read tracker entries for which a full cache line read has missed the L3. The GQ read tracker L3 miss to fill occupancy count is divided by this count to obtain the average cache line read L3 miss latency. The latency represents the time after which the L3 has determined that the cache line has missed. The time between a GQ read tracker allocation and the L3 determining that the cache line has missed is the average L3 hit latency. The total L3 cache line read miss latency is the hit latency + L3 miss latency. |         |
| 03H        | 04H         | UNC_GQ_ALLOC.RT_TO_L3_RESP                 | Counts the number of GQ read tracker entries that are allocated in the read tracker queue that hit or miss the L3. The GQ read tracker L3 hit occupancy count is divided by this count to obtain the average L3 hit latency.   |         |
| 03H        | 08H         | UNC_GQ_ALLOC.RT_TO_RTID_ACQUIRED           | Counts the number of GQ read tracker entries that are allocated in the read tracker, have missed in the L3 and have not acquired a Request Transaction ID. The GQ read tracker L3 miss to RTID acquired occupancy count is divided by this count to obtain the average latency for a read L3 miss to acquire an RTID.  |         |
| 03H        | 10H         | UNC_GQ_ALLOC.WT_TO_RTID_ACQUIRED           | Counts the number of GQ write tracker entries that are allocated in the write tracker, have missed in the L3 and have not acquired a Request Transaction ID. The GQ write tracker L3 miss to RTID occupancy count is divided by this count to obtain the average latency for a write L3 miss to acquire an RTID.   |         |
| 03H        | 20H         | UNC_GQ_ALLOC.WRITE_TRACKER                 | Counts the number of GQ write tracker entries that are allocated in the write tracker queue that miss the L3. The GQ write tracker occupancy count is divided by this count to obtain the average L3 write miss latency.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                    | Description  | Comment |
|------------|-------------|--|--|---------|
| 03H        | 40H         | UNC_GQ_ALLOC.PEER_PROBE_TRACKER        | Counts the number of GQ peer probe tracker (snoop) entries that are allocated in the peer probe tracker queue that miss the L3. The GQ peer probe occupancy count is divided by this count to obtain the average L3 peer probe miss latency. |         |
| 04H        | 01H         | UNC_GQ_DATA.FROM_QPI                   | Cycles Global Queue Quickpath Interface input data port is busy importing data from the Quickpath Interface. Each cycle the input port can transfer 8 or 16 bytes of data.   |         |
| 04H        | 02H         | UNC_GQ_DATA.FROM_QMC                   | Cycles Global Queue Quickpath Memory Interface input data port is busy importing data from the Quickpath Memory Interface. Each cycle the input port can transfer 8 or 16 bytes of data.   |         |
| 04H        | 04H         | UNC_GQ_DATA.FROM_L3                    | Cycles GQ L3 input data port is busy importing data from the Last Level Cache. Each cycle the input port can transfer 32 bytes of data.  |         |
| 04H        | 08H         | UNC_GQ_DATA.FROM_CORES_02              | Cycles GQ Core 0 and 2 input data port is busy importing data from processor cores 0 and 2. Each cycle the input port can transfer 32 bytes of data.   |         |
| 04H        | 10H         | UNC_GQ_DATA.FROM_CORES_13              | Cycles GQ Core 1 and 3 input data port is busy importing data from processor cores 1 and 3. Each cycle the input port can transfer 32 bytes of data.   |         |
| 05H        | 01H         | UNC_GQ_DATA.TO_QPI_QMC                 | Cycles GQ QPI and QMC output data port is busy sending data to the Quickpath Interface or Quickpath Memory Interface. Each cycle the output port can transfer 32 bytes of data.  |         |
| 05H        | 02H         | UNC_GQ_DATA.TO_L3                      | Cycles GQ L3 output data port is busy sending data to the Last Level Cache. Each cycle the output port can transfer 32 bytes of data.  |         |
| 05H        | 04H         | UNC_GQ_DATA.TO_CORES                   | Cycles GQ Core output data port is busy sending data to the Cores. Each cycle the output port can transfer 32 bytes of data.   |         |
| 06H        | 01H         | UNC_SNP_RESP_TO_LOCAL_HOME.I_STATE     | Number of snoop responses to the local home that L3 does not have the referenced cache line.   |         |
| 06H        | 02H         | UNC_SNP_RESP_TO_LOCAL_HOME.S_STATE     | Number of snoop responses to the local home that L3 has the referenced line cached in the S state.   |         |
| 06H        | 04H         | UNC_SNP_RESP_TO_LOCAL_HOME.FWD_S_STATE | Number of responses to code or data read snoops to the local home that the L3 has the referenced cache line in the E state. The L3 cache line state is changed to the S state and the line is forwarded to the local home in the S state.    |         |
| 06H        | 08H         | UNC_SNP_RESP_TO_LOCAL_HOME.FWD_I_STATE | Number of responses to read invalidate snoops to the local home that the L3 has the referenced cache line in the M state. The L3 cache line state is invalidated and the line is forwarded to the local home in the M state.                 |         |
| 06H        | 10H         | UNC_SNP_RESP_TO_LOCAL_HOME.CONFLICT    | Number of conflict snoop responses sent to the local home.   |         |

**Table 19-20. Performance Events In the Processor Uncore for  
Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment |
|------------|-------------|---|---|---------|
| 06H        | 20H         | UNC_SNP_RESP_TO_LOCAL_HOME.WB           | Number of responses to code or data read snoops to the local home that the L3 has the referenced line cached in the M state.  |         |
| 07H        | 01H         | UNC_SNP_RESP_TO_REMOTE_HOME.I_STATE     | Number of snoop responses to a remote home that L3 does not have the referenced cache line.   |         |
| 07H        | 02H         | UNC_SNP_RESP_TO_REMOTE_HOME.S_STATE     | Number of snoop responses to a remote home that L3 has the referenced line cached in the S state.   |         |
| 07H        | 04H         | UNC_SNP_RESP_TO_REMOTE_HOME.FWD_S_STATE | Number of responses to code or data read snoops to a remote home that the L3 has the referenced cache line in the E state. The L3 cache line state is changed to the S state and the line is forwarded to the remote home in the S state. |         |
| 07H        | 08H         | UNC_SNP_RESP_TO_REMOTE_HOME.FWD_I_STATE | Number of responses to read invalidate snoops to a remote home that the L3 has the referenced cache line in the M state. The L3 cache line state is invalidated and the line is forwarded to the remote home in the M state.              |         |
| 07H        | 10H         | UNC_SNP_RESP_TO_REMOTE_HOME.CONFLICT    | Number of conflict snoop responses sent to the local home.  |         |
| 07H        | 20H         | UNC_SNP_RESP_TO_REMOTE_HOME.WB          | Number of responses to code or data read snoops to a remote home that the L3 has the referenced line cached in the M state.   |         |
| 07H        | 24H         | UNC_SNP_RESP_TO_REMOTE_HOME.HITM        | Number of HITM snoop responses to a remote home.  |         |
| 08H        | 01H         | UNC_L3_HITS.READ                        | Number of code read, data read and RFO requests that hit in the L3.   |         |
| 08H        | 02H         | UNC_L3_HITS.WRITE                       | Number of writeback requests that hit in the L3. Writebacks from the cores will always result in L3 hits due to the inclusive property of the L3.   |         |
| 08H        | 04H         | UNC_L3_HITS.PROBE                       | Number of snoops from IOH or remote sockets that hit in the L3.   |         |
| 08H        | 03H         | UNC_L3_HITS.ANY                         | Number of reads and writes that hit the L3.   |         |
| 09H        | 01H         | UNC_L3_MISS.READ                        | Number of code read, data read and RFO requests that miss the L3.   |         |
| 09H        | 02H         | UNC_L3_MISS.WRITE                       | Number of writeback requests that miss the L3. Should always be zero as writebacks from the cores will always result in L3 hits due to the inclusive property of the L3.  |         |
| 09H        | 04H         | UNC_L3_MISS.PROBE                       | Number of snoops from IOH or remote sockets that miss the L3.   |         |
| 09H        | 03H         | UNC_L3_MISS.ANY                         | Number of reads and writes that miss the L3.  |         |
| 0AH        | 01H         | UNC_L3_LINES_IN.M_STATE                 | Counts the number of L3 lines allocated in M state. The only time a cache line is allocated in the M state is when the line was forwarded in M state is forwarded due to a Snoop Read Invalidate Own request.                             |         |
| 0AH        | 02H         | UNC_L3_LINES_IN.E_STATE                 | Counts the number of L3 lines allocated in E state.   |         |
| 0AH        | 04H         | UNC_L3_LINES_IN.S_STATE                 | Counts the number of L3 lines allocated in S state.   |         |



**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description   | Comment |
|------------|-------------|---------------------------------|---|---------|
| 0AH        | 08H         | UNC_L3_LINES_IN.F_STATE         | Counts the number of L3 lines allocated in F state.   |         |
| 0AH        | 0FH         | UNC_L3_LINES_IN.ANY             | Counts the number of L3 lines allocated in any state.   |         |
| 0BH        | 01H         | UNC_L3_LINES_OUT.M_STATE        | Counts the number of L3 lines victimized that were in the M state. When the victim cache line is in M state, the line is written to its home cache agent which can be either local or remote. |         |
| 0BH        | 02H         | UNC_L3_LINES_OUT.E_STATE        | Counts the number of L3 lines victimized that were in the E state.  |         |
| 0BH        | 04H         | UNC_L3_LINES_OUT.S_STATE        | Counts the number of L3 lines victimized that were in the S state.  |         |
| 0BH        | 08H         | UNC_L3_LINES_OUT.I_STATE        | Counts the number of L3 lines victimized that were in the I state.  |         |
| 0BH        | 10H         | UNC_L3_LINES_OUT.F_STATE        | Counts the number of L3 lines victimized that were in the F state.  |         |
| 0BH        | 1FH         | UNC_L3_LINES_OUT.ANY            | Counts the number of L3 lines victimized in any state.  |         |
| 20H        | 01H         | UNC_QHL_REQUESTS.IOH_READS      | Counts number of Quickpath Home Logic read requests from the IOH.   |         |
| 20H        | 02H         | UNC_QHL_REQUESTS.IOH_WRITES     | Counts number of Quickpath Home Logic write requests from the IOH.  |         |
| 20H        | 04H         | UNC_QHL_REQUESTS.REMOTE_READS   | Counts number of Quickpath Home Logic read requests from a remote socket.   |         |
| 20H        | 08H         | UNC_QHL_REQUESTS.REMOTE_WRITES  | Counts number of Quickpath Home Logic write requests from a remote socket.  |         |
| 20H        | 10H         | UNC_QHL_REQUESTS.LOCAL_READS    | Counts number of Quickpath Home Logic read requests from the local socket.  |         |
| 20H        | 20H         | UNC_QHL_REQUESTS.LOCAL_WRITES   | Counts number of Quickpath Home Logic write requests from the local socket.   |         |
| 21H        | 01H         | UNC_QHL_CYCLES_FULL.IOH         | Counts uclk cycles all entries in the Quickpath Home Logic IOH are full.  |         |
| 21H        | 02H         | UNC_QHL_CYCLES_FULL.REMOTE      | Counts uclk cycles all entries in the Quickpath Home Logic remote tracker are full.   |         |
| 21H        | 04H         | UNC_QHL_CYCLES_FULL.LOCAL       | Counts uclk cycles all entries in the Quickpath Home Logic local tracker are full.  |         |
| 22H        | 01H         | UNC_QHL_CYCLES_NOT_EMPTY.IOH    | Counts uclk cycles all entries in the Quickpath Home Logic IOH is busy.   |         |
| 22H        | 02H         | UNC_QHL_CYCLES_NOT_EMPTY.REMOTE | Counts uclk cycles all entries in the Quickpath Home Logic remote tracker is busy.  |         |
| 22H        | 04H         | UNC_QHL_CYCLES_NOT_EMPTY.LOCAL  | Counts uclk cycles all entries in the Quickpath Home Logic local tracker is busy.   |         |
| 23H        | 01H         | UNC_QHL_OCCUPANCY.IOH           | QHL IOH tracker allocate to deallocate read occupancy.  |         |
| 23H        | 02H         | UNC_QHL_OCCUPANCY.REMOTE        | QHL remote tracker allocate to deallocate read occupancy.   |         |
| 23H        | 04H         | UNC_QHL_OCCUPANCY.LOCAL         | QHL local tracker allocate to deallocate read occupancy.  |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description  | Comment |
|------------|-------------|--------------------------------|--|---------|
| 24H        | 02H         | UNC_QHL_ADDRESS_CONFLICTS.2WAY | Counts number of QHL Active Address Table (AAT) entries that saw a max of 2 conflicts. The AAT is a structure that tracks requests that are in conflict. The requests themselves are in the home tracker entries. The count is reported when an AAT entry deallocates. |         |
| 24H        | 04H         | UNC_QHL_ADDRESS_CONFLICTS.3WAY | Counts number of QHL Active Address Table (AAT) entries that saw a max of 3 conflicts. The AAT is a structure that tracks requests that are in conflict. The requests themselves are in the home tracker entries. The count is reported when an AAT entry deallocates. |         |
| 25H        | 01H         | UNC_QHL_CONFLICT_CYCLES.IOH    | Counts cycles the Quickpath Home Logic IOH Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.   |         |
| 25H        | 02H         | UNC_QHL_CONFLICT_CYCLES.REMOTE | Counts cycles the Quickpath Home Logic Remote Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.  |         |
| 25H        | 04H         | UNC_QHL_CONFLICT_CYCLES.LOCAL  | Counts cycles the Quickpath Home Logic Local Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.   |         |
| 26H        | 01H         | UNC_QHL_TO_QMC_BYPASS          | Counts number of requests to the Quickpath Memory Controller that bypass the Quickpath Home Logic. All local accesses can be bypassed. For remote requests, only read requests can be bypassed.  |         |
| 27H        | 01H         | UNC_QMC_NORMAL_FULL.READ.CH0   | Uncore cycles all the entries in the DRAM channel 0 medium or low priority queue are occupied with read requests.  |         |
| 27H        | 02H         | UNC_QMC_NORMAL_FULL.READ.CH1   | Uncore cycles all the entries in the DRAM channel 1 medium or low priority queue are occupied with read requests.  |         |
| 27H        | 04H         | UNC_QMC_NORMAL_FULL.READ.CH2   | Uncore cycles all the entries in the DRAM channel 2 medium or low priority queue are occupied with read requests.  |         |
| 27H        | 08H         | UNC_QMC_NORMAL_FULL.WRITE.CH0  | Uncore cycles all the entries in the DRAM channel 0 medium or low priority queue are occupied with write requests.   |         |
| 27H        | 10H         | UNC_QMC_NORMAL_FULL.WRITE.CH1  | Counts cycles all the entries in the DRAM channel 1 medium or low priority queue are occupied with write requests.   |         |
| 27H        | 20H         | UNC_QMC_NORMAL_FULL.WRITE.CH2  | Uncore cycles all the entries in the DRAM channel 2 medium or low priority queue are occupied with write requests.   |         |
| 28H        | 01H         | UNC_QMC_ISOC_FULL.READ.CH0     | Counts cycles all the entries in the DRAM channel 0 high priority queue are occupied with isochronous read requests.   |         |
| 28H        | 02H         | UNC_QMC_ISOC_FULL.READ.CH1     | Counts cycles all the entries in the DRAM channel 1 high priority queue are occupied with isochronous read requests.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic          | Description  | Comment |
|------------|-------------|------------------------------|--|---------|
| 28H        | 04H         | UNC_QMC_ISOC_FULL.READ.C H2  | Counts cycles all the entries in the DRAM channel 2 high priority queue are occupied with isochronous read requests.   |         |
| 28H        | 08H         | UNC_QMC_ISOC_FULL.WRITE.C H0 | Counts cycles all the entries in the DRAM channel 0 high priority queue are occupied with isochronous write requests.  |         |
| 28H        | 10H         | UNC_QMC_ISOC_FULL.WRITE.C H1 | Counts cycles all the entries in the DRAM channel 1 high priority queue are occupied with isochronous write requests.  |         |
| 28H        | 20H         | UNC_QMC_ISOC_FULL.WRITE.C H2 | Counts cycles all the entries in the DRAM channel 2 high priority queue are occupied with isochronous write requests.  |         |
| 29H        | 01H         | UNC_QMC_BUSY.READ.CH0        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 0.   |         |
| 29H        | 02H         | UNC_QMC_BUSY.READ.CH1        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 1.   |         |
| 29H        | 04H         | UNC_QMC_BUSY.READ.CH2        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 2.   |         |
| 29H        | 08H         | UNC_QMC_BUSY.WRITE.CH0       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 0.  |         |
| 29H        | 10H         | UNC_QMC_BUSY.WRITE.CH1       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 1.  |         |
| 29H        | 20H         | UNC_QMC_BUSY.WRITE.CH2       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 2.  |         |
| 2AH        | 01H         | UNC_QMC_OCCUPANCY.CH0        | IMC channel 0 normal read request occupancy.   |         |
| 2AH        | 02H         | UNC_QMC_OCCUPANCY.CH1        | IMC channel 1 normal read request occupancy.   |         |
| 2AH        | 04H         | UNC_QMC_OCCUPANCY.CH2        | IMC channel 2 normal read request occupancy.   |         |
| 2BH        | 01H         | UNC_QMC_ISSOC_OCCUPANCY.CH0  | IMC channel 0 issoc read request occupancy.  |         |
| 2BH        | 02H         | UNC_QMC_ISSOC_OCCUPANCY.CH1  | IMC channel 1 issoc read request occupancy.  |         |
| 2BH        | 04H         | UNC_QMC_ISSOC_OCCUPANCY.CH2  | IMC channel 2 issoc read request occupancy.  |         |
| 2BH        | 07H         | UNC_QMC_ISSOC_READS.ANY      | IMC issoc read request occupancy.  |         |
| 2CH        | 01H         | UNC_QMC_NORMAL_READS.C H0    | Counts the number of Quickpath Memory Controller channel 0 medium and low priority read requests. The QMC channel 0 normal read occupancy divided by this count provides the average QMC channel 0 read latency. |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description  | Comment |
|------------|-------------|-------------------------------------|--|---------|
| 2CH        | 02H         | UNC_QMC_NORMAL_READS.CH1            | Counts the number of Quickpath Memory Controller channel 1 medium and low priority read requests. The QMC channel 1 normal read occupancy divided by this count provides the average QMC channel 1 read latency. |         |
| 2CH        | 04H         | UNC_QMC_NORMAL_READS.CH2            | Counts the number of Quickpath Memory Controller channel 2 medium and low priority read requests. The QMC channel 2 normal read occupancy divided by this count provides the average QMC channel 2 read latency. |         |
| 2CH        | 07H         | UNC_QMC_NORMAL_READS.ANY            | Counts the number of Quickpath Memory Controller medium and low priority read requests. The QMC normal read occupancy divided by this count provides the average QMC read latency.                               |         |
| 2DH        | 01H         | UNC_QMC_HIGH_PRIORITY_READS.CH0     | Counts the number of Quickpath Memory Controller channel 0 high priority isochronous read requests.  |         |
| 2DH        | 02H         | UNC_QMC_HIGH_PRIORITY_READS.CH1     | Counts the number of Quickpath Memory Controller channel 1 high priority isochronous read requests.  |         |
| 2DH        | 04H         | UNC_QMC_HIGH_PRIORITY_READS.CH2     | Counts the number of Quickpath Memory Controller channel 2 high priority isochronous read requests.  |         |
| 2DH        | 07H         | UNC_QMC_HIGH_PRIORITY_READS.ANY     | Counts the number of Quickpath Memory Controller high priority isochronous read requests.  |         |
| 2EH        | 01H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH0 | Counts the number of Quickpath Memory Controller channel 0 critical priority isochronous read requests.  |         |
| 2EH        | 02H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH1 | Counts the number of Quickpath Memory Controller channel 1 critical priority isochronous read requests.  |         |
| 2EH        | 04H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH2 | Counts the number of Quickpath Memory Controller channel 2 critical priority isochronous read requests.  |         |
| 2EH        | 07H         | UNC_QMC_CRITICAL_PRIORITY_READS.ANY | Counts the number of Quickpath Memory Controller critical priority isochronous read requests.  |         |
| 2FH        | 01H         | UNC_QMC_WRITES.FULL.CH0             | Counts number of full cache line writes to DRAM channel 0.   |         |
| 2FH        | 02H         | UNC_QMC_WRITES.FULL.CH1             | Counts number of full cache line writes to DRAM channel 1.   |         |
| 2FH        | 04H         | UNC_QMC_WRITES.FULL.CH2             | Counts number of full cache line writes to DRAM channel 2.   |         |
| 2FH        | 07H         | UNC_QMC_WRITES.FULL.ANY             | Counts number of full cache line writes to DRAM.   |         |
| 2FH        | 08H         | UNC_QMC_WRITES.PARTIAL.CH0          | Counts number of partial cache line writes to DRAM channel 0.  |         |
| 2FH        | 10H         | UNC_QMC_WRITES.PARTIAL.CH1          | Counts number of partial cache line writes to DRAM channel 1.  |         |
| 2FH        | 20H         | UNC_QMC_WRITES.PARTIAL.CH2          | Counts number of partial cache line writes to DRAM channel 2.  |         |
| 2FH        | 38H         | UNC_QMC_WRITES.PARTIAL.ANY          | Counts number of partial cache line writes to DRAM.  |         |
| 30H        | 01H         | UNC_QMC_CANCEL.CH0                  | Counts number of DRAM channel 0 cancel requests.   |         |
| 30H        | 02H         | UNC_QMC_CANCEL.CH1                  | Counts number of DRAM channel 1 cancel requests.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                          | Description   | Comment |
|------------|-------------|--|---|---------|
| 30H        | 04H         | UNC_QMC_CANCEL.CH2                           | Counts number of DRAM channel 2 cancel requests.  |         |
| 30H        | 07H         | UNC_QMC_CANCEL.ANY                           | Counts number of DRAM cancel requests.  |         |
| 31H        | 01H         | UNC_QMC_PRIORITY_UPDATE S.CH0                | Counts number of DRAM channel 0 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |
| 31H        | 02H         | UNC_QMC_PRIORITY_UPDATE S.CH1                | Counts number of DRAM channel 1 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |
| 31H        | 04H         | UNC_QMC_PRIORITY_UPDATE S.CH2                | Counts number of DRAM channel 2 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |
| 31H        | 07H         | UNC_QMC_PRIORITY_UPDATE S.ANY                | Counts number of DRAM priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request.           |         |
| 33H        | 04H         | UNC_QHL_FRC_ACK_CNFLTS.L OCAL                | Counts number of Force Acknowledge Conflict messages sent by the Quickpath Home Logic to the local home.  |         |
| 40H        | 01H         | UNC_QPI_TX_STALLED_SINGL E_FLIT.HOME.LINK_0  | Counts cycles the Quickpath outbound link 0 HOME virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |
| 40H        | 02H         | UNC_QPI_TX_STALLED_SINGL E_FLIT.SNOOP.LINK_0 | Counts cycles the Quickpath outbound link 0 SNOOP virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.  |         |
| 40H        | 04H         | UNC_QPI_TX_STALLED_SINGL E_FLIT.NDR.LINK_0   | Counts cycles the Quickpath outbound link 0 non-data response virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                                      |         |
| 40H        | 08H         | UNC_QPI_TX_STALLED_SINGL E_FLIT.HOME.LINK_1  | Counts cycles the Quickpath outbound link 1 HOME virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description   | Comment |
|------------|-------------|---|---|---------|
| 40H        | 10H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.SNOOP.LINK_1 | Counts cycles the Quickpath outbound link 1 SNOOP virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                |         |
| 40H        | 20H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.NDR.LINK_1   | Counts cycles the Quickpath outbound link 1 non-data response virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.    |         |
| 40H        | 07H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.LINK_0       | Counts cycles the Quickpath outbound link 0 virtual channels are stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                    |         |
| 40H        | 38H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.LINK_1       | Counts cycles the Quickpath outbound link 1 virtual channels are stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                    |         |
| 41H        | 01H         | UNC_QPI_TX_STALLED_MULTIFLIT.DRS.LINK_0     | Counts cycles the Quickpath outbound link 0 Data Response virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.         |         |
| 41H        | 02H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCB.LINK_0     | Counts cycles the Quickpath outbound link 0 Non-Coherent Bypass virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |
| 41H        | 04H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCS.LINK_0     | Counts cycles the Quickpath outbound link 0 Non-Coherent Standard virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated. |         |
| 41H        | 08H         | UNC_QPI_TX_STALLED_MULTIFLIT.DRS.LINK_1     | Counts cycles the Quickpath outbound link 1 Data Response virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.         |         |
| 41H        | 10H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCB.LINK_1     | Counts cycles the Quickpath outbound link 1 Non-Coherent Bypass virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment |
|------------|-------------|---|---|---------|
| 41H        | 20H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCS.LINK_1 | Counts cycles the Quickpath outbound link 1 Non-Coherent Standard virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated. |         |
| 41H        | 07H         | UNC_QPI_TX_STALLED_MULTIFLIT.LINK_0     | Counts cycles the Quickpath outbound link 0 virtual channels are stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                     |         |
| 41H        | 38H         | UNC_QPI_TX_STALLED_MULTIFLIT.LINK_1     | Counts cycles the Quickpath outbound link 1 virtual channels are stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                     |         |
| 42H        | 02H         | UNC_QPI_TX_HEADER.BUSY.LINK_0           | Number of cycles that the header buffer in the Quickpath Interface outbound link 0 is busy.   |         |
| 42H        | 08H         | UNC_QPI_TX_HEADER.BUSY.LINK_1           | Number of cycles that the header buffer in the Quickpath Interface outbound link 1 is busy.   |         |
| 43H        | 01H         | UNC_QPI_RX_NO_PPT_CREDIT.STALLS.LINK_0  | Number of cycles that snoop packets incoming to the Quickpath Interface link 0 are stalled and not sent to the GQ because the GQ Peer Probe Tracker (PPT) does not have any available entries.  |         |
| 43H        | 02H         | UNC_QPI_RX_NO_PPT_CREDIT.STALLS.LINK_1  | Number of cycles that snoop packets incoming to the Quickpath Interface link 1 are stalled and not sent to the GQ because the GQ Peer Probe Tracker (PPT) does not have any available entries.  |         |
| 60H        | 01H         | UNC_DRAM_OPEN.CH0                       | Counts number of DRAM Channel 0 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |
| 60H        | 02H         | UNC_DRAM_OPEN.CH1                       | Counts number of DRAM Channel 1 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |
| 60H        | 04H         | UNC_DRAM_OPEN.CH2                       | Counts number of DRAM Channel 2 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |
| 61H        | 01H         | UNC_DRAM_PAGE_CLOSE.CH0                 | DRAM channel 0 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |
| 61H        | 02H         | UNC_DRAM_PAGE_CLOSE.CH1                 | DRAM channel 1 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |
| 61H        | 04H         | UNC_DRAM_PAGE_CLOSE.CH2                 | DRAM channel 2 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |

**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description   | Comment |
|------------|-------------|---------------------------------|---|---------|
| 62H        | 01H         | UNC_DRAM_PAGE_MISS.CHO          | Counts the number of precharges (PRE) that were issued to DRAM channel 0 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 62H        | 02H         | UNC_DRAM_PAGE_MISS.CH1          | Counts the number of precharges (PRE) that were issued to DRAM channel 1 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 62H        | 04H         | UNC_DRAM_PAGE_MISS.CH2          | Counts the number of precharges (PRE) that were issued to DRAM channel 2 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 63H        | 01H         | UNC_DRAM_READ_CAS.CHO           | Counts the number of times a read CAS command was issued on DRAM channel 0.   |         |
| 63H        | 02H         | UNC_DRAM_READ_CAS.AUTO PRE_CH0  | Counts the number of times a read CAS command was issued on DRAM channel 0 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 63H        | 04H         | UNC_DRAM_READ_CAS.CH1           | Counts the number of times a read CAS command was issued on DRAM channel 1.   |         |
| 63H        | 08H         | UNC_DRAM_READ_CAS.AUTO PRE_CH1  | Counts the number of times a read CAS command was issued on DRAM channel 1 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 63H        | 10H         | UNC_DRAM_READ_CAS.CH2           | Counts the number of times a read CAS command was issued on DRAM channel 2.   |         |
| 63H        | 20H         | UNC_DRAM_READ_CAS.AUTO PRE_CH2  | Counts the number of times a read CAS command was issued on DRAM channel 2 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 64H        | 01H         | UNC_DRAM_WRITE_CAS.CHO          | Counts the number of times a write CAS command was issued on DRAM channel 0.  |         |
| 64H        | 02H         | UNC_DRAM_WRITE_CAS.AUTO PRE_CH0 | Counts the number of times a write CAS command was issued on DRAM channel 0 where the command issued used the auto-precharge (auto page close) mode.  |         |
| 64H        | 04H         | UNC_DRAM_WRITE_CAS.CH1          | Counts the number of times a write CAS command was issued on DRAM channel 1.  |         |
| 64H        | 08H         | UNC_DRAM_WRITE_CAS.AUTO PRE_CH1 | Counts the number of times a write CAS command was issued on DRAM channel 1 where the command issued used the auto-precharge (auto page close) mode.  |         |
| 64H        | 10H         | UNC_DRAM_WRITE_CAS.CH2          | Counts the number of times a write CAS command was issued on DRAM channel 2.  |         |



**Table 19-20. Performance Events In the Processor Uncore for Intel® Core™ i7 Processor and Intel® Xeon® Processor 5500 Series (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                | Description  | Comment |
|------------|-------------|------------------------------------|--|---------|
| 64H        | 20H         | UNC_DRAM_WRITE_CAS.AUTO<br>PRE_CH2 | Counts the number of times a write CAS command was issued on DRAM channel 2 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 65H        | 01H         | UNC_DRAM_REFRESH.CH0               | Counts number of DRAM channel 0 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 65H        | 02H         | UNC_DRAM_REFRESH.CH1               | Counts number of DRAM channel 1 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 65H        | 04H         | UNC_DRAM_REFRESH.CH2               | Counts number of DRAM channel 2 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 66H        | 01H         | UNC_DRAM_PRE_ALL.CH0               | Counts number of DRAM Channel 0 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |
| 66H        | 02H         | UNC_DRAM_PRE_ALL.CH1               | Counts number of DRAM Channel 1 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |
| 66H        | 04H         | UNC_DRAM_PRE_ALL.CH2               | Counts number of DRAM Channel 2 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |

Intel Xeon processors with CUID signature of DisplayFamily\_DisplayModel 06\_2EH have a distinct uncore sub-system that is significantly different from the uncore found in processors with CUID signature 06\_1AH, 06\_1EH, and 06\_1FH. Model-specific performance monitoring events for its uncore will be available in future documentation.

## 19.10 PERFORMANCE MONITORING EVENTS FOR PROCESSORS BASED ON INTEL® MICROARCHITECTURE CODE NAME WESTMERE

Intel 64 processors based on Intel® microarchitecture code name Westmere support the architectural and model-specific performance monitoring events listed in Table 19-1 and Table 19-21. Table 19-21 applies to processors with CUID signature of DisplayFamily\_DisplayModel encoding with the following values: 06\_25H, 06\_2CH. In addition, these processors (CUID signature of DisplayFamily\_DisplayModel 06\_25H, 06\_2CH) also support the following model-specific, product-specific uncore performance monitoring events listed in Table 19-22. Fixed counters support the architecture events defined in Table 19-2.

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere**

| Event Num. | Umask Value | Event Mask Mnemonic      | Description                                    | Comment |
|------------|-------------|--------------------------|--|---------|
| 03H        | 02H         | LOAD_BLOCK.OVERLAP_STORE | Loads that partially overlap an earlier store. |         |
| 04H        | 07H         | SB_DRAIN.ANY             | All Store buffer stall cycles.                 |         |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description   | Comment                              |
|------------|-------------|--|---|--------------------------------------|
| 05H        | 02H         | MISALIGN_MEMORY.STORE                    | All store referenced with misaligned address.   |                                      |
| 06H        | 04H         | STORE_BLOCKS.AT_RET                      | Counts number of loads delayed with at-Retirement block code. The following loads need to be executed at retirement and wait for all senior stores on the same thread to be drained: load splitting across 4K boundary (page split), load accessing uncacheable (UC or WC) memory, load lock, and load with page table in UC or WC memory region. |                                      |
| 06H        | 08H         | STORE_BLOCKS.L1D_BLOCK                   | Cacheable loads delayed with L1D block code.  |                                      |
| 07H        | 01H         | PARTIAL_ADDRESS_ALIAS                    | Counts false dependency due to partial address aliasing.  |                                      |
| 08H        | 01H         | DTLB_LOAD_MISSES.ANY                     | Counts all load misses that cause a page walk.  |                                      |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED          | Counts number of completed page walks due to load miss in the STLB.   |                                      |
| 08H        | 04H         | DTLB_LOAD_MISSES.WALK_CYCLES             | Cycles PMH is busy with a page walk due to a load miss in the STLB.   |                                      |
| 08H        | 10H         | DTLB_LOAD_MISSES.STLB_HIT                | Number of cache load STLB hits.   |                                      |
| 08H        | 20H         | DTLB_LOAD_MISSES.PDE_MISSES              | Number of DTLB cache load misses where the low part of the linear to physical address translation was missed.   |                                      |
| 0BH        | 01H         | MEM_INST_RETIRED.LOADS                   | Counts the number of instructions with an architecturally-visible load retired on the architected path.   |                                      |
| 0BH        | 02H         | MEM_INST_RETIRED.STORES                  | Counts the number of instructions with an architecturally-visible store retired on the architected path.  |                                      |
| 0BH        | 10H         | MEM_INST_RETIRED.LATENCY_ABOVE_THRESHOLD | Counts the number of instructions exceeding the latency specified with Id_lat facility.   | In conjunction with Id_lat facility. |
| 0CH        | 01H         | MEM_STORE_RETIRED.DTLB_MISS              | The event counts the number of retired stores that missed the DTLB. The DTLB miss is not counted if the store operation causes a fault. Does not counter prefetches. Counts both primary and secondary misses to the TLB.   |                                      |
| 0EH        | 01H         | UOPS_ISSUED.ANY                          | Counts the number of Uops issued by the Register Allocation Table to the Reservation Station, i.e. the UOPs issued from the front end to the back end.  |                                      |
| 0EH        | 01H         | UOPS_ISSUED.STALLED_CYCLES               | Counts the number of cycles no uops issued by the Register Allocation Table to the Reservation Station, i.e. the UOPs issued from the front end to the back end.  | Set "invert=1, cmask = 1".           |
| 0EH        | 02H         | UOPS_ISSUED.FUSED                        | Counts the number of fused Uops that were issued from the Register Allocation Table to the Reservation Station.   |                                      |
| 0FH        | 01H         | MEM_UNCORE_RETIRED.UNKNOWNSOURCE         | Load instructions retired with unknown LLC miss (Precise Event).  | Applicable to one and two sockets.   |
| 0FH        | 02H         | MEM_UNCORE_RETIRED.OHTE_R_CORE_L2_HIT    | Load instructions retired that HIT modified data in sibling core (Precise Event).   | Applicable to one and two sockets.   |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                                | Description   | Comment                            |
|------------|-------------|--|---|------------------------------------|
| 0FH        | 04H         | MEM_UNCORE_RETIRED.REMOTE_HITM                     | Load instructions retired that HIT modified data in remote socket (Precise Event).  | Applicable to two sockets only.    |
| 0FH        | 08H         | MEM_UNCORE_RETIRED.LOCAL_DRAM_AND_REMOTE_CACHE_HIT | Load instructions retired local dram and remote cache HIT data sources (Precise Event).   | Applicable to one and two sockets. |
| 0FH        | 10H         | MEM_UNCORE_RETIRED.REMOTE_DRAM                     | Load instructions retired remote DRAM and remote home-remote cache HITM (Precise Event).  | Applicable to two sockets only.    |
| 0FH        | 20H         | MEM_UNCORE_RETIRED.OTHER_LLC_MISS                  | Load instructions retired other LLC miss (Precise Event).   | Applicable to two sockets only.    |
| 0FH        | 80H         | MEM_UNCORE_RETIRED.UNCACHEABLE                     | Load instructions retired I/O (Precise Event).  | Applicable to one and two sockets. |
| 10H        | 01H         | FP_COMP_OPS_EXE.X87                                | Counts the number of FP Computational Uops Executed. The number of FADD, FSUB, FCOM, FMULs, integer MULs and IMULs, FDIVs, FPREMs, FSQRTS, integer DIVs, and IDIVs. This event does not distinguish an FADD used in the middle of a transcendental flow from a separate FADD instruction. |                                    |
| 10H        | 02H         | FP_COMP_OPS_EXE.MMX                                | Counts number of MMX Uops executed.   |                                    |
| 10H        | 04H         | FP_COMP_OPS_EXE.SSE_FP                             | Counts number of SSE and SSE2 FP uops executed.   |                                    |
| 10H        | 08H         | FP_COMP_OPS_EXE.SSE2_INTEGER                       | Counts number of SSE2 integer uops executed.  |                                    |
| 10H        | 10H         | FP_COMP_OPS_EXE.SSE_FP_PACKED                      | Counts number of SSE FP packed uops executed.   |                                    |
| 10H        | 20H         | FP_COMP_OPS_EXE.SSE_FP_SCALAR                      | Counts number of SSE FP scalar uops executed.   |                                    |
| 10H        | 40H         | FP_COMP_OPS_EXE.SSE_SINGLE_PRECISION               | Counts number of SSE* FP single precision uops executed.  |                                    |
| 10H        | 80H         | FP_COMP_OPS_EXE.SSE_DOUBLE_PRECISION               | Counts number of SSE* FP double precision uops executed.  |                                    |
| 12H        | 01H         | SIMD_INT_128.PACKED_MPY                            | Counts number of 128 bit SIMD integer multiply operations.  |                                    |
| 12H        | 02H         | SIMD_INT_128.PACKED_SHIFT                          | Counts number of 128 bit SIMD integer shift operations.   |                                    |
| 12H        | 04H         | SIMD_INT_128.PACK                                  | Counts number of 128 bit SIMD integer pack operations.  |                                    |
| 12H        | 08H         | SIMD_INT_128.UNPACK                                | Counts number of 128 bit SIMD integer unpack operations.  |                                    |
| 12H        | 10H         | SIMD_INT_128.PACKED_LOGICAL                        | Counts number of 128 bit SIMD integer logical operations.   |                                    |
| 12H        | 20H         | SIMD_INT_128.PACKED_ARITH                          | Counts number of 128 bit SIMD integer arithmetic operations.  |                                    |
| 12H        | 40H         | SIMD_INT_128.SHUFFLE_MOVE                          | Counts number of 128 bit SIMD integer shuffle and move operations.  |                                    |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic      | Description  | Comment  |
|------------|-------------|--------------------------|--|--|
| 13H        | 01H         | LOAD_DISPATCH.RS         | Counts number of loads dispatched from the Reservation Station that bypass the Memory Order Buffer.  |  |
| 13H        | 02H         | LOAD_DISPATCH.RS_DELAYED | Counts the number of delayed RS dispatches at the stage latch. If an RS dispatch cannot bypass to LB, it has another chance to dispatch from the one-cycle delayed staging latch before it is written into the LB.   |  |
| 13H        | 04H         | LOAD_DISPATCH.MOB        | Counts the number of loads dispatched from the Reservation Station to the Memory Order Buffer.   |  |
| 13H        | 07H         | LOAD_DISPATCH.ANY        | Counts all loads dispatched from the Reservation Station.  |  |
| 14H        | 01H         | ARITH.CYCLES_DIV_BUSY    | Counts the number of cycles the divider is busy executing divide or square root operations. The divide can be integer, X87 or Streaming SIMD Extensions (SSE). The square root operation can be either X87 or SSE. Set 'edge = 1, invert = 1, cmask = 1' to count the number of divides.   | Count may be incorrect when SMT is on.   |
| 14H        | 02H         | ARITH.MUL                | Counts the number of multiply operations executed. This includes integer as well as floating point multiply operations but excludes DPPS mul and MPSAD.  | Count may be incorrect when SMT is on.   |
| 17H        | 01H         | INST_QUEUE_WRITES        | Counts the number of instructions written into the instruction queue every cycle.  |  |
| 18H        | 01H         | INST_DECODED.DECO        | Counts number of instructions that require decoder 0 to be decoded. Usually, this means that the instruction maps to more than 1 uop.  |  |
| 19H        | 01H         | TWO_UOP_INSTS_DECODED    | An instruction that generates two uops was decoded.  |  |
| 1EH        | 01H         | INST_QUEUE_WRITE_CYCLES  | This event counts the number of cycles during which instructions are written to the instruction queue. Dividing this counter by the number of instructions written to the instruction queue (INST_QUEUE_WRITES) yields the average number of instructions decoded each cycle. If this number is less than four and the pipe stalls, this indicates that the decoder is failing to decode enough instructions per cycle to sustain the 4-wide pipeline. | If SSE* instructions that are 6 bytes or longer arrive one after another, then front end throughput may limit execution speed. |
| 20H        | 01H         | LSD_OVERFLOW             | Number of loops that cannot stream from the instruction queue.   |  |
| 24H        | 01H         | L2_RQSTS.LD_HIT          | Counts number of loads that hit the L2 cache. L2 loads include both L1D demand misses as well as L1D prefetches. L2 loads can be rejected for various reasons. Only non rejected loads are counted.  |  |
| 24H        | 02H         | L2_RQSTS.LD_MISS         | Counts the number of loads that miss the L2 cache. L2 loads include both L1D demand misses as well as L1D prefetches.  |  |
| 24H        | 03H         | L2_RQSTS.LOADS           | Counts all L2 load requests. L2 loads include both L1D demand misses as well as L1D prefetches.  |  |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic          | Description   | Comment |
|------------|-------------|------------------------------|---|---------|
| 24H        | 04H         | L2_RQSTS.RFO_HIT             | Counts the number of store RFO requests that hit the L2 cache. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches. Count includes WC memory requests, where the data is not fetched but the permission to write the line is required. |         |
| 24H        | 08H         | L2_RQSTS.RFO_MISS            | Counts the number of store RFO requests that miss the L2 cache. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches.   |         |
| 24H        | 0CH         | L2_RQSTS.RFOS                | Counts all L2 store RFO requests. L2 RFO requests include both L1D demand RFO misses as well as L1D RFO prefetches.   |         |
| 24H        | 10H         | L2_RQSTS.IFETCH_HIT          | Counts number of instruction fetches that hit the L2 cache. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.  |         |
| 24H        | 20H         | L2_RQSTS.IFETCH_MISS         | Counts number of instruction fetches that miss the L2 cache. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.   |         |
| 24H        | 30H         | L2_RQSTS.IFETCHES            | Counts all instruction fetches. L2 instruction fetches include both L1I demand misses as well as L1I instruction prefetches.  |         |
| 24H        | 40H         | L2_RQSTS.PREFETCH_HIT        | Counts L2 prefetch hits for both code and data.   |         |
| 24H        | 80H         | L2_RQSTS.PREFETCH_MISS       | Counts L2 prefetch misses for both code and data.   |         |
| 24H        | C0H         | L2_RQSTS.PREFETCHES          | Counts all L2 prefetches for both code and data.  |         |
| 24H        | AAH         | L2_RQSTS.MISS                | Counts all L2 misses for both code and data.  |         |
| 24H        | FFH         | L2_RQSTS.REFERENCES          | Counts all L2 requests for both code and data.  |         |
| 26H        | 01H         | L2_DATA_RQSTS.DEMAND.I_STATE | Counts number of L2 data demand loads where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss. L2 demand loads are both L1D demand misses and L1D prefetches.   |         |
| 26H        | 02H         | L2_DATA_RQSTS.DEMAND.S_STATE | Counts number of L2 data demand loads where the cache line to be loaded is in the S (shared) state. L2 demand loads are both L1D demand misses and L1D prefetches.  |         |
| 26H        | 04H         | L2_DATA_RQSTS.DEMAND.E_STATE | Counts number of L2 data demand loads where the cache line to be loaded is in the E (exclusive) state. L2 demand loads are both L1D demand misses and L1D prefetches.   |         |
| 26H        | 08H         | L2_DATA_RQSTS.DEMAND.M_STATE | Counts number of L2 data demand loads where the cache line to be loaded is in the M (modified) state. L2 demand loads are both L1D demand misses and L1D prefetches.  |         |
| 26H        | 0FH         | L2_DATA_RQSTS.DEMAND.MESI    | Counts all L2 data demand requests. L2 demand loads are both L1D demand misses and L1D prefetches.  |         |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment                       |
|------------|-------------|--------------------------------|---|-------------------------------|
| 26H        | 10H         | L2_DATA_RQSTS.PREFETCH.I_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss.  |                               |
| 26H        | 20H         | L2_DATA_RQSTS.PREFETCH.S_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the S (shared) state. A prefetch RFO will miss on an S state line, while a prefetch read will hit on an S state line. |                               |
| 26H        | 40H         | L2_DATA_RQSTS.PREFETCH.E_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the E (exclusive) state.  |                               |
| 26H        | 80H         | L2_DATA_RQSTS.PREFETCH.M_STATE | Counts number of L2 prefetch data loads where the cache line to be loaded is in the M (modified) state.   |                               |
| 26H        | F0H         | L2_DATA_RQSTS.PREFETCH.MESI    | Counts all L2 prefetch requests.  |                               |
| 26H        | FFH         | L2_DATA_RQSTS.ANY              | Counts all L2 data requests.  |                               |
| 27H        | 01H         | L2_WRITE.RFO.I_STATE           | Counts number of L2 demand store RFO requests where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss. The L1D prefetcher does not issue a RFO prefetch.                    | This is a demand RFO request. |
| 27H        | 02H         | L2_WRITE.RFO.S_STATE           | Counts number of L2 store RFO requests where the cache line to be loaded is in the S (shared) state. The L1D prefetcher does not issue a RFO prefetch.  | This is a demand RFO request. |
| 27H        | 08H         | L2_WRITE.RFO.M_STATE           | Counts number of L2 store RFO requests where the cache line to be loaded is in the M (modified) state. The L1D prefetcher does not issue a RFO prefetch.  | This is a demand RFO request. |
| 27H        | 0EH         | L2_WRITE.RFO.HIT               | Counts number of L2 store RFO requests where the cache line to be loaded is in either the S, E or M states. The L1D prefetcher does not issue a RFO prefetch.   | This is a demand RFO request. |
| 27H        | 0FH         | L2_WRITE.RFO.MESI              | Counts all L2 store RFO requests. The L1D prefetcher does not issue a RFO prefetch.   | This is a demand RFO request. |
| 27H        | 10H         | L2_WRITE.LOCK.I_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the I (invalid) state, i.e., a cache miss.   |                               |
| 27H        | 20H         | L2_WRITE.LOCK.S_STATE          | Counts number of L2 lock RFO requests where the cache line to be loaded is in the S (shared) state.   |                               |
| 27H        | 40H         | L2_WRITE.LOCK.E_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the E (exclusive) state.   |                               |
| 27H        | 80H         | L2_WRITE.LOCK.M_STATE          | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in the M (modified) state.  |                               |
| 27H        | E0H         | L2_WRITE.LOCK.HIT              | Counts number of L2 demand lock RFO requests where the cache line to be loaded is in either the S, E, or M state.   |                               |
| 27H        | F0H         | L2_WRITE.LOCK.MESI             | Counts all L2 demand lock RFO requests.   |                               |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description  | Comment            |
|------------|-------------|----------------------------------|--|--------------------|
| 28H        | 01H         | L1D_WB_L2.I_STATE                | Counts number of L1 writebacks to the L2 where the cache line to be written is in the I (invalid) state, i.e., a cache miss.   |                    |
| 28H        | 02H         | L1D_WB_L2.S_STATE                | Counts number of L1 writebacks to the L2 where the cache line to be written is in the S state.   |                    |
| 28H        | 04H         | L1D_WB_L2.E_STATE                | Counts number of L1 writebacks to the L2 where the cache line to be written is in the E (exclusive) state.   |                    |
| 28H        | 08H         | L1D_WB_L2.M_STATE                | Counts number of L1 writebacks to the L2 where the cache line to be written is in the M (modified) state.  |                    |
| 28H        | 0FH         | L1D_WB_L2.MESI                   | Counts all L1 writebacks to the L2 .   |                    |
| 2EH        | 41H         | L3_LAT_CACHE.MISS                | Counts uncore Last Level Cache misses. Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended.                                  | See Table 19-1.    |
| 2EH        | 4FH         | L3_LAT_CACHE.REFERENCE           | Counts uncore Last Level Cache references. Because cache hierarchy, cache sizes and other implementation-specific characteristics; value comparison to estimate performance differences is not recommended.                              | See Table 19-1.    |
| 3CH        | 00H         | CPU_CLK_UNHALTED.THREAD_P        | Counts the number of thread cycles while the thread is not in a halt state. The thread enters the halt state when it is running the HLT instruction. The core frequency may change from time to time due to power or thermal throttling. | See Table 19-1.    |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF_P           | Increments at the frequency of TSC when not halted.  | See Table 19-1.    |
| 49H        | 01H         | DTLB_MISSES.ANY                  | Counts the number of misses in the STLB which causes a page walk.  |                    |
| 49H        | 02H         | DTLB_MISSES.WALK_COMPLETED       | Counts number of misses in the STLB which resulted in a completed page walk.   |                    |
| 49H        | 04H         | DTLB_MISSES.WALK_CYCLES          | Counts cycles of page walk due to misses in the STLB.  |                    |
| 49H        | 10H         | DTLB_MISSES.STLB_HIT             | Counts the number of DTLB first level misses that hit in the second level TLB. This event is only relevant if the core contains multiple DTLB levels.  |                    |
| 49H        | 20H         | DTLB_MISSES.PDE_MISS             | Number of DTLB misses caused by low part of address, includes references to 2M pages because 2M pages do not use the PDE.  |                    |
| 49H        | 80H         | DTLB_MISSES.LARGE_WALK_COMPLETED | Counts number of completed large page walks due to misses in the STLB.   |                    |
| 4CH        | 01H         | LOAD_HIT_PRE                     | Counts load operations sent to the L1 data cache while a previous SSE prefetch instruction to the same cache line has started prefetching but has not yet finished.  | Counter 0, 1 only. |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                           | Description   | Comment   |
|------------|-------------|---|---|---|
| 4EH        | 01H         | L1D_PREFETCH.REQUESTS                         | Counts number of hardware prefetch requests dispatched out of the prefetch FIFO.  | Counter 0, 1 only.  |
| 4EH        | 02H         | L1D_PREFETCH.MISS                             | Counts number of hardware prefetch requests that miss the L1D. There are two prefetchers in the L1D. A streamer, which predicts lines sequentially after this one should be fetched, and the IP prefetcher that remembers access patterns for the current instruction. The streamer prefetcher stops on an L1D hit, while the IP prefetcher does not. | Counter 0, 1 only.  |
| 4EH        | 04H         | L1D_PREFETCH.TRIGGERS                         | Counts number of prefetch requests triggered by the Finite State Machine and pushed into the prefetch FIFO. Some of the prefetch requests are dropped due to overwrites or competition between the IP index prefetcher and streamer prefetcher. The prefetch FIFO contains 4 entries.   | Counter 0, 1 only.  |
| 4FH        | 10H         | EPT.WALK_CYCLES                               | Counts Extended Page walk cycles.   |   |
| 51H        | 01H         | L1D.REPL                                      | Counts the number of lines brought into the L1 data cache.  | Counter 0, 1 only.  |
| 51H        | 02H         | L1D.M_REPL                                    | Counts the number of modified lines brought into the L1 data cache.   | Counter 0, 1 only.  |
| 51H        | 04H         | L1D.M_EVICT                                   | Counts the number of modified lines evicted from the L1 data cache due to replacement.  | Counter 0, 1 only.  |
| 51H        | 08H         | L1D.M_SNOOP_EVICT                             | Counts the number of modified lines evicted from the L1 data cache due to snoop HITM intervention.  | Counter 0, 1 only.  |
| 52H        | 01H         | L1D_CACHE_PREFETCH_LOCK_FB_HIT                | Counts the number of cacheable load lock speculated instructions accepted into the fill buffer.   |   |
| 60H        | 01H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND.READ_DATA | Counts weighted cycles of offcore demand data read requests. Does not include L2 prefetch requests.   | Counter 0.  |
| 60H        | 02H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND.READ_CODE | Counts weighted cycles of offcore demand code read requests. Does not include L2 prefetch requests.   | Counter 0.  |
| 60H        | 04H         | OFFCORE_REQUESTS_OUTSTANDING.DEMAND.RFO       | Counts weighted cycles of offcore demand RFO requests. Does not include L2 prefetch requests.   | Counter 0.  |
| 60H        | 08H         | OFFCORE_REQUESTS_OUTSTANDING.ANY.READ         | Counts weighted cycles of offcore read requests of any kind. Include L2 prefetch requests.  | Counter 0.  |
| 63H        | 01H         | CACHE_LOCK_CYCLES.L1D_L2                      | Cycle count during which the L1D and L2 are locked. A lock is asserted when there is a locked memory access, due to uncacheable memory, a locked operation that spans two cache lines, or a page walk from an uncacheable page table. This event does not cause locks, it merely detects them.  | Counter 0, 1 only. L1D and L2 locks have a very high performance penalty and it is highly recommended to avoid such accesses. |
| 63H        | 02H         | CACHE_LOCK_CYCLES.L1D                         | Counts the number of cycles that cacheline in the L1 data cache unit is locked.   | Counter 0, 1 only.  |
| 6CH        | 01H         | IO_TRANSACTIONS                               | Counts the number of completed I/O transactions.  |   |
| 80H        | 01H         | L1I.HITS                                      | Counts all instruction fetches that hit the L1 instruction cache.   |   |



**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description  | Comment |
|------------|-------------|----------------------------------|--|---------|
| 80H        | 02H         | L1I.MISSES                       | Counts all instruction fetches that miss the L1I cache. This includes instruction cache misses, streaming buffer misses, victim cache misses and uncacheable fetches. An instruction fetch miss is counted only once and not once for every cycle it is outstanding. |         |
| 80H        | 03H         | L1I.READS                        | Counts all instruction fetches, including uncacheable fetches that bypass the L1I.   |         |
| 80H        | 04H         | L1I.CYCLES_STALLED               | Cycle counts for which an instruction fetch stalls due to a L1I cache miss, ITLB miss or ITLB fault.   |         |
| 82H        | 01H         | LARGE_ITLB.HIT                   | Counts number of large ITLB hits.  |         |
| 85H        | 01H         | ITLB_MISSES.ANY                  | Counts the number of misses in all levels of the ITLB which causes a page walk.  |         |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETED       | Counts number of misses in all levels of the ITLB which resulted in a completed page walk.   |         |
| 85H        | 04H         | ITLB_MISSES.WALK_CYCLES          | Counts ITLB miss page walk cycles.   |         |
| 85H        | 10H         | ITLB_MISSES.STLB_HIT             | Counts number of ITLB first level miss but second level hits.  |         |
| 85H        | 80H         | ITLB_MISSES.LARGE_WALK_COMPLETED | Counts number of completed large page walks due to misses in the STLB.   |         |
| 87H        | 01H         | ILD_STALL.LCP                    | Cycles Instruction Length Decoder stalls due to length changing prefixes: 66, 67 or REX.W (for Intel 64) instructions which change the length of the decoded instruction.  |         |
| 87H        | 02H         | ILD_STALL.MRU                    | Instruction Length Decoder stall cycles due to Branch Prediction Unit (PBU) Most Recently Used (MRU) bypass.   |         |
| 87H        | 04H         | ILD_STALL.IQ_FULL                | Stall cycles due to a full instruction queue.  |         |
| 87H        | 08H         | ILD_STALL.REGEN                  | Counts the number of regen stalls.   |         |
| 87H        | 0FH         | ILD_STALL.ANY                    | Counts any cycles the Instruction Length Decoder is stalled.   |         |
| 88H        | 01H         | BR_INST_EXEC.COND                | Counts the number of conditional near branch instructions executed, but not necessarily retired.   |         |
| 88H        | 02H         | BR_INST_EXEC.DIRECT              | Counts all unconditional near branch instructions excluding calls and indirect branches.   |         |
| 88H        | 04H         | BR_INST_EXEC.INDIRECT_NON_CALL   | Counts the number of executed indirect near branch instructions that are not calls.  |         |
| 88H        | 07H         | BR_INST_EXEC.NON_CALLS           | Counts all non-call near branch instructions executed, but not necessarily retired.  |         |
| 88H        | 08H         | BR_INST_EXEC.RETURN_NEAR         | Counts indirect near branches that have a return mnemonic.   |         |
| 88H        | 10H         | BR_INST_EXEC.DIRECT_NEAR_CALL    | Counts unconditional near call branch instructions, excluding non-call branch, executed.   |         |
| 88H        | 20H         | BR_INST_EXEC.INDIRECT_NEAR_CALL  | Counts indirect near calls, including both register and memory indirect, executed.   |         |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description   | Comment  |
|------------|-------------|---------------------------------|---|--|
| 88H        | 30H         | BR_INST_EXEC.NEAR_CALLS         | Counts all near call branches executed, but not necessarily retired.  |  |
| 88H        | 40H         | BR_INST_EXEC.TAKEN              | Counts taken near branches executed, but not necessarily retired.   |  |
| 88H        | 7FH         | BR_INST_EXEC.ANY                | Counts all near executed branches (not necessarily retired). This includes only instructions and not micro-op branches. Frequent branching is not necessarily a major performance issue. However frequent branch mispredictions may be a problem.   |  |
| 89H        | 01H         | BR_MISP_EXEC.COND               | Counts the number of mispredicted conditional near branch instructions executed, but not necessarily retired.   |  |
| 89H        | 02H         | BR_MISP_EXEC.DIRECT             | Counts mispredicted macro unconditional near branch instructions, excluding calls and indirect branches (should always be 0).   |  |
| 89H        | 04H         | BR_MISP_EXEC.INDIRECT_NO_N_CALL | Counts the number of executed mispredicted indirect near branch instructions that are not calls.  |  |
| 89H        | 07H         | BR_MISP_EXEC.NON_CALLS          | Counts mispredicted non-call near branches executed, but not necessarily retired.   |  |
| 89H        | 08H         | BR_MISP_EXEC.RETURN_NEAR        | Counts mispredicted indirect branches that have a rear return mnemonic.   |  |
| 89H        | 10H         | BR_MISP_EXEC.DIRECT_NEAR_CALL   | Counts mispredicted non-indirect near calls executed, (should always be 0).   |  |
| 89H        | 20H         | BR_MISP_EXEC.INDIRECT_NEAR_CALL | Counts mispredicted indirect near calls executed, including both register and memory indirect.  |  |
| 89H        | 30H         | BR_MISP_EXEC.NEAR_CALLS         | Counts all mispredicted near call branches executed, but not necessarily retired.   |  |
| 89H        | 40H         | BR_MISP_EXEC.TAKEN              | Counts executed mispredicted near branches that are taken, but not necessarily retired.   |  |
| 89H        | 7FH         | BR_MISP_EXEC.ANY                | Counts the number of mispredicted near branch instructions that were executed, but not necessarily retired.   |  |
| A2H        | 01H         | RESOURCE_STALLS.ANY             | Counts the number of Allocator resource related stalls. Includes register renaming buffer entries, memory buffer entries. In addition to resource related stalls, this event counts some other events. Includes stalls arising during branch misprediction recovery, such as if retirement of the mispredicted branch is delayed and stalls arising while store buffer is draining from synchronizing operations. | Does not include stalls due to SuperQ (off core) queue full, too many cache misses, etc. |
| A2H        | 02H         | RESOURCE_STALLS.LOAD            | Counts the cycles of stall due to lack of load buffer for load operation.   |  |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description  | Comment   |
|------------|-------------|-----------------------------------|--|---|
| A2H        | 04H         | RESOURCE_STALLS.RS_FULL           | This event counts the number of cycles when the number of instructions in the pipeline waiting for execution reaches the limit the processor can handle. A high count of this event indicates that there are long latency operations in the pipe (possibly load and store operations that miss the L2 cache, or instructions dependent upon instructions further down the pipeline that have yet to retire.  | When RS is full, new instructions cannot enter the reservation station and start execution. |
| A2H        | 08H         | RESOURCE_STALLS.STORE             | This event counts the number of cycles that a resource related stall will occur due to the number of store instructions reaching the limit of the pipeline, (i.e. all store buffers are used). The stall ends when a store instruction commits its data to the cache or memory.  |   |
| A2H        | 10H         | RESOURCE_STALLS.ROB_FULL          | Counts the cycles of stall due to re-order buffer full.  |   |
| A2H        | 20H         | RESOURCE_STALLS.FPCW              | Counts the number of cycles while execution was stalled due to writing the floating-point unit (FPU) control word.   |   |
| A2H        | 40H         | RESOURCE_STALLS.MXCSR             | Stalls due to the MXCSR register rename occurring to close to a previous MXCSR rename. The MXCSR provides control and status for the MMX registers.  |   |
| A2H        | 80H         | RESOURCE_STALLS.OTHER             | Counts the number of cycles while execution was stalled due to other resource issues.  |   |
| A6H        | 01H         | MACRO_INSTS.FUSIONS_DECODED       | Counts the number of instructions decoded that are macro-fused but not necessarily executed or retired.  |   |
| A7H        | 01H         | BACLEAR_FORCE_IQ                  | Counts number of times a BACLEAR was forced by the Instruction Queue. The IQ is also responsible for providing conditional branch prediction direction based on a static scheme and dynamic data provided by the L2 Branch Prediction Unit. If the conditional branch target is not found in the Target Array and the IQ predicts that the branch is taken, then the IQ will force the Branch Address Calculator to issue a BACLEAR. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline. |   |
| A8H        | 01H         | LSD.UOPS                          | Counts the number of micro-ops delivered by loop stream detector.  | Use cmask=1 and invert to count cycles.   |
| AEH        | 01H         | ITLB_FLUSH                        | Counts the number of ITLB flushes.   |   |
| B0H        | 01H         | OFFCORE_REQUESTS.DEMAND.READ_DATA | Counts number of offcore demand data read requests. Does not count L2 prefetch requests.   |   |
| B0H        | 02H         | OFFCORE_REQUESTS.DEMAND.READ_CODE | Counts number of offcore demand code read requests. Does not count L2 prefetch requests.   |   |
| B0H        | 04H         | OFFCORE_REQUESTS.DEMAND.RFO       | Counts number of offcore demand RFO requests. Does not count L2 prefetch requests.   |   |
| B0H        | 08H         | OFFCORE_REQUESTS.ANY.READ         | Counts number of offcore read requests. Includes L2 prefetch requests.   |   |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description  | Comment                                      |
|------------|-------------|---|--|--|
| B0H        | 10H         | OFFCORE_REQUESTS.ANY.RFO                  | Counts number of offcore RFO requests. Includes L2 prefetch requests.  |  |
| B0H        | 40H         | OFFCORE_REQUESTS.L1D_WRITEBACK            | Counts number of L1D writebacks to the uncore.   |  |
| B0H        | 80H         | OFFCORE_REQUESTS.ANY                      | Counts all offcore requests.   |  |
| B1H        | 01H         | UOPS_EXECUTED.PORT0                       | Counts number of uops executed that were issued on port 0. Port 0 handles integer arithmetic, SIMD and FP add uops.  |  |
| B1H        | 02H         | UOPS_EXECUTED.PORT1                       | Counts number of uops executed that were issued on port 1. Port 1 handles integer arithmetic, SIMD, integer shift, FP multiply and FP divide uops.   |  |
| B1H        | 04H         | UOPS_EXECUTED.PORT2_CORE                  | Counts number of uops executed that were issued on port 2. Port 2 handles the load uops. This is a core count only and cannot be collected per thread.   |  |
| B1H        | 08H         | UOPS_EXECUTED.PORT3_CORE                  | Counts number of uops executed that were issued on port 3. Port 3 handles store uops. This is a core count only and cannot be collected per thread.  |  |
| B1H        | 10H         | UOPS_EXECUTED.PORT4_CORE                  | Counts number of uops executed that where issued on port 4. Port 4 handles the value to be stored for the store uops issued on port 3. This is a core count only and cannot be collected per thread. |  |
| B1H        | 1FH         | UOPS_EXECUTED.CORE_ACTIVE_CYCLES_NO_PORTS | Counts number of cycles there are one or more uops being executed and were issued on ports 0-4. This is a core count only and cannot be collected per thread.  |  |
| B1H        | 20H         | UOPS_EXECUTED.PORT5                       | Counts number of uops executed that where issued on port 5.  |  |
| B1H        | 3FH         | UOPS_EXECUTED.CORE_ACTIVE_CYCLES          | Counts number of cycles there are one or more uops being executed on any ports. This is a core count only and cannot be collected per thread.  |  |
| B1H        | 40H         | UOPS_EXECUTED.PORT015                     | Counts number of uops executed that where issued on port 0, 1, or 5.   | Use cmask=1, invert=1 to count stall cycles. |
| B1H        | 80H         | UOPS_EXECUTED.PORT234                     | Counts number of uops executed that where issued on port 2, 3, or 4.   |  |
| B2H        | 01H         | OFFCORE_REQUESTS_SQ_FULL                  | Counts number of cycles the SQ is full to handle off-core requests.  |  |
| B3H        | 01H         | SNOOPQ_REQUESTS_OUTSTANDING.DATA          | Counts weighted cycles of snoopq requests for data. Counter 0 only.  | Use cmask=1 to count cycles not empty.       |
| B3H        | 02H         | SNOOPQ_REQUESTS_OUTSTANDING.INVALIDATE    | Counts weighted cycles of snoopq invalidate requests. Counter 0 only.  | Use cmask=1 to count cycles not empty.       |
| B3H        | 04H         | SNOOPQ_REQUESTS_OUTSTANDING.CODE          | Counts weighted cycles of snoopq requests for code. Counter 0 only.  | Use cmask=1 to count cycles not empty.       |
| B4H        | 01H         | SNOOPQ_REQUESTS.CODE                      | Counts the number of snoop code requests.  |  |
| B4H        | 02H         | SNOOPQ_REQUESTS.DATA                      | Counts the number of snoop data requests.  |  |
| B4H        | 04H         | SNOOPQ_REQUESTS.INVALIDATE                | Counts the number of snoop invalidate requests.  |  |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic          | Description   | Comment  |
|------------|-------------|------------------------------|---|--|
| B7H        | 01H         | OFF_CORE_RESPONSE_0          | See Section 18.3.1.1.3, "Off-core Response Performance Monitoring in the Processor Core".   | Requires programming MSR 01A6H.  |
| B8H        | 01H         | SNOOP_RESPONSE.HIT           | Counts HIT snoop response sent by this thread in response to a snoop request.   |  |
| B8H        | 02H         | SNOOP_RESPONSE.HITE          | Counts HIT E snoop response sent by this thread in response to a snoop request.   |  |
| B8H        | 04H         | SNOOP_RESPONSE.HITM          | Counts HIT M snoop response sent by this thread in response to a snoop request.   |  |
| BBH        | 01H         | OFF_CORE_RESPONSE_1          | See Section 18.3.1.1.3, "Off-core Response Performance Monitoring in the Processor Core".   | Use MSR 01A7H.   |
| C0H        | 00H         | INST_RETIRED.ANY_P           | See Table 19-1.<br>Notes: INST_RETIRED.ANY is counted by a designated fixed counter. INST_RETIRED.ANY_P is counted by a programmable counter and is an architectural performance event. Event is supported if CPUID.A.EBX[1] = 0.   | Counting: Faulting executions of GETSEC/VM entry/VM Exit/MWait will not count as retired instructions. |
| C0H        | 02H         | INST_RETIRED.X87             | Counts the number of floating point computational operations retired: floating point computational operations executed by the assist handler and sub-operations of complex floating point instructions like transcendental instructions.  |  |
| C0H        | 04H         | INST_RETIRED.MMX             | Counts the number of retired: MMX instructions.   |  |
| C2H        | 01H         | UOPS_RETIRED.ANY             | Counts the number of micro-ops retired, (macro-fused=1, micro-fused=2, others=1; maximum count of 8 per cycle). Most instructions are composed of one or two micro-ops. Some instructions are decoded into longer sequences such as repeat instructions, floating point transcendental instructions, and assists. | Use cmask=1 and invert to count active cycles or stalled cycles.                                       |
| C2H        | 02H         | UOPS_RETIRED.RETIRE_SLOT     | Counts the number of retirement slots used each cycle.  |  |
| C2H        | 04H         | UOPS_RETIRED.MACRO_FUSED     | Counts number of macro-fused uops retired.  |  |
| C3H        | 01H         | MACHINE_CLEAR.CYCLES         | Counts the cycles machine clear is asserted.  |  |
| C3H        | 02H         | MACHINE_CLEAR.MEM_ORDER      | Counts the number of machine clears due to memory order conflicts.  |  |
| C3H        | 04H         | MACHINE_CLEAR.SMC            | Counts the number of times that a program writes to a code section. Self-modifying code causes a severe penalty in all Intel 64 and IA-32 processors. The modified cache line is written back to the L2 and L3 caches.  |  |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES | Branch instructions at retirement.  | See Table 19-1.  |
| C4H        | 01H         | BR_INST_RETIRED.CONDITIONAL  | Counts the number of conditional branch instructions retired.   |  |
| C4H        | 02H         | BR_INST_RETIRED.NEAR_CALL    | Counts the number of direct & indirect near unconditional calls retired.  |  |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                          | Description   | Comment         |
|------------|-------------|--|---|-----------------|
| C5H        | 00H         | BR_MISP_RETIRE.D.ALL_BRAN<br>CHES            | Mispredicted branch instructions at retirement.   | See Table 19-1. |
| C5H        | 01H         | BR_MISP_RETIRE.D.CONDITION<br>AL             | Counts mispredicted conditional retired calls.  |                 |
| C5H        | 02H         | BR_MISP_RETIRE.D.NEAR_CAL<br>L               | Counts mispredicted direct & indirect near unconditional retired calls.   |                 |
| C5H        | 04H         | BR_MISP_RETIRE.D.ALL_BRAN<br>CHES            | Counts all mispredicted retired calls.  |                 |
| C7H        | 01H         | SSEX_UOPS_RETIRE.D.PACKED<br>_SINGLE         | Counts SIMD packed single-precision floating-point uops retired.  |                 |
| C7H        | 02H         | SSEX_UOPS_RETIRE.D.SCALAR<br>_SINGLE         | Counts SIMD scalar single-precision floating-point uops retired.  |                 |
| C7H        | 04H         | SSEX_UOPS_RETIRE.D.PACKED<br>_DOUBLE         | Counts SIMD packed double-precision floating-point uops retired.  |                 |
| C7H        | 08H         | SSEX_UOPS_RETIRE.D.SCALAR<br>_DOUBLE         | Counts SIMD scalar double-precision floating-point uops retired.  |                 |
| C7H        | 10H         | SSEX_UOPS_RETIRE.D.VECTOR<br>_INTEGER        | Counts 128-bit SIMD vector integer uops retired.  |                 |
| C8H        | 20H         | ITLB_MISS_RETIRE.D                           | Counts the number of retired instructions that missed the ITLB when the instruction was fetched.  |                 |
| CBH        | 01H         | MEM_LOAD_RETIRE.D.L1D_HIT                    | Counts number of retired loads that hit the L1 data cache.  |                 |
| CBH        | 02H         | MEM_LOAD_RETIRE.D.L2_HIT                     | Counts number of retired loads that hit the L2 data cache.  |                 |
| CBH        | 04H         | MEM_LOAD_RETIRE.D.L3_UNSH<br>ARED_HIT        | Counts number of retired loads that hit their own, unshared lines in the L3 cache.  |                 |
| CBH        | 08H         | MEM_LOAD_RETIRE.D.OTHER_<br>CORE_L2_HIT_HITM | Counts number of retired loads that hit in a sibling core's L2 (on die core). Since the L3 is inclusive of all cores on the package, this is an L3 hit. This counts both clean and modified hits.   |                 |
| CBH        | 10H         | MEM_LOAD_RETIRE.D.L3_MISS                    | Counts number of retired loads that miss the L3 cache. The load was satisfied by a remote socket, local memory or an IOH.   |                 |
| CBH        | 40H         | MEM_LOAD_RETIRE.D.HIT_LFB                    | Counts number of retired loads that miss the L1D and the address is located in an allocated line fill buffer and will soon be committed to cache. This is counting secondary L1D misses.  |                 |
| CBH        | 80H         | MEM_LOAD_RETIRE.D.DTLB_MI<br>SS              | Counts the number of retired loads that missed the DTLB. The DTLB miss is not counted if the load operation causes a fault. This event counts loads from cacheable memory only. The event does not count loads by software prefetches. Counts both primary and secondary misses to the TLB. |                 |
| CCH        | 01H         | FP_MMX_TRANS.TO_FP                           | Counts the first floating-point instruction following any MMX instruction. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |                 |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic        | Description   | Comment |
|------------|-------------|----------------------------|---|---------|
| CCH        | 02H         | FP_MMX_TRANS.TO_MMX        | Counts the first MMX instruction following a floating-point instruction. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |         |
| CCH        | 03H         | FP_MMX_TRANS.ANY           | Counts all transitions from floating point to MMX instructions and from MMX instructions to floating point instructions. You can use this event to estimate the penalties for the transitions between floating-point and MMX technology states.   |         |
| DOH        | 01H         | MACRO_INSTS.DECODED        | Counts the number of instructions decoded, (but not necessarily executed or retired).   |         |
| D1H        | 01H         | UOPS_DECODED.STALL_CYCLE S | Counts the cycles of decoder stalls. INV=1, Cmask=1.  |         |
| D1H        | 02H         | UOPS_DECODED.MS            | Counts the number of Uops decoded by the Microcode Sequencer, MS. The MS delivers uops when the instruction is more than 4 uops long or a microcode assist is occurring.  |         |
| D1H        | 04H         | UOPS_DECODED.ESP_FOLDIN G  | Counts number of stack pointer (ESP) instructions decoded: push, pop, call, ret, etc. ESP instructions do not generate a Uop to increment or decrement ESP. Instead, they update an ESP_Offset register that keeps track of the delta to the current value of the ESP register.   |         |
| D1H        | 08H         | UOPS_DECODED.ESP_SYNC      | Counts number of stack pointer (ESP) sync operations where an ESP instruction is corrected by adding the ESP offset register to the current value of the ESP register.  |         |
| D2H        | 01H         | RAT_STALLS.FLAGS           | Counts the number of cycles during which execution stalled due to several reasons, one of which is a partial flag register stall. A partial register stall may occur when two conditions are met: 1) an instruction modifies some, but not all, of the flags in the flag register and 2) the next instruction, which depends on flags, depends on flags that were not modified by this instruction.                 |         |
| D2H        | 02H         | RAT_STALLS.REGISTERS       | This event counts the number of cycles instruction execution latency became longer than the defined latency because the instruction used a register that was partially written by previous instruction.   |         |
| D2H        | 04H         | RAT_STALLS.ROB_READ_POR T  | Counts the number of cycles when ROB read port stalls occurred, which did not allow new micro-ops to enter the out-of-order pipeline. Note that, at this stage in the pipeline, additional stalls may occur at the same cycle and prevent the stalled micro-ops from entering the pipe. In such a case, micro-ops retry entering the execution pipe in the next cycle and the ROB-read port stall is counted again. |         |
| D2H        | 08H         | RAT_STALLS.SCOREBOARD      | Counts the cycles where we stall due to microarchitecturally required serialization. Microcode scoreboarding stalls.  |         |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic  | Description  | Comment   |
|------------|-------------|----------------------|--|---|
| D2H        | 0FH         | RAT_STALLS.ANY       | Counts all Register Allocation Table stall cycles due to: Cycles when ROB read port stalls occurred, which did not allow new micro-ops to enter the execution pipe, Cycles when partial register stalls occurred, Cycles when flag stalls occurred, Cycles floating-point unit (FPU) status word stalls occurred. To count each of these conditions separately use the events: RAT_STALLS.ROB_READ_PORT, RAT_STALLS.PARTIAL, RAT_STALLS.FLAGS, and RAT_STALLS.FPSW.                      |   |
| D4H        | 01H         | SEG_RENAME_STALLS    | Counts the number of stall cycles due to the lack of renaming resources for the ES, DS, FS, and GS segment registers. If a segment is renamed but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.   |   |
| D5H        | 01H         | ES_REG_RENAMES       | Counts the number of times the ES segment register is renamed.   |   |
| DBH        | 01H         | UOP_UNFUSION         | Counts unfusion events due to floating point exception to a fused uop.   |   |
| E0H        | 01H         | BR_INST_DECODED      | Counts the number of branch instructions decoded.  |   |
| E5H        | 01H         | BPU_MISSED_CALL_RET  | Counts number of times the Branch Prediction Unit missed predicting a call or return branch.   |   |
| E6H        | 01H         | BACLEAR.CLEAR        | Counts the number of times the front end is resteeered, mainly when the Branch Prediction Unit cannot provide a correct prediction and this is corrected by the Branch Address Calculator at the front end. This can occur if the code has many branches such that they cannot be consumed by the BPU. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline. The effect on total execution time depends on the surrounding code. |   |
| E6H        | 02H         | BACLEAR.BAD_TARGET   | Counts number of Branch Address Calculator clears (BACLEAR) asserted due to conditional branch instructions in which there was a target hit but the direction was wrong. Each BACLEAR asserted by the BAC generates approximately an 8 cycle bubble in the instruction fetch pipeline.   |   |
| E8H        | 01H         | BPU_CLEARS.EARLY     | Counts early (normal) Branch Prediction Unit clears: BPU predicted a taken branch after incorrectly assuming that it was not taken.  | The BPU clear leads to 2 cycle bubble in the front end. |
| E8H        | 02H         | BPU_CLEARS.LATE      | Counts late Branch Prediction Unit clears due to Most Recently Used conflicts. The PBU clear leads to a 3 cycle bubble in the front end.   |   |
| ECH        | 01H         | THREAD_ACTIVE        | Counts cycles threads are active.  |   |
| F0H        | 01H         | L2_TRANSACTIONS.LOAD | Counts L2 load operations due to HW prefetch or demand loads.  |   |



**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic         | Description  | Comment |
|------------|-------------|-----------------------------|--|---------|
| F0H        | 02H         | L2_TRANSACTION.S.RFO        | Counts L2 RFO operations due to HW prefetch or demand RFOs.  |         |
| F0H        | 04H         | L2_TRANSACTION.S.IFETCH     | Counts L2 instruction fetch operations due to HW prefetch or demand ifetch.  |         |
| F0H        | 08H         | L2_TRANSACTION.S.PREFETCH   | Counts L2 prefetch operations.   |         |
| F0H        | 10H         | L2_TRANSACTION.S.L1D_WB     | Counts L1D writeback operations to the L2.   |         |
| F0H        | 20H         | L2_TRANSACTION.S.FILL       | Counts L2 cache line fill operations due to load, RFO, L1D writeback or prefetch.  |         |
| F0H        | 40H         | L2_TRANSACTION.S.WB         | Counts L2 writeback operations to the L3.  |         |
| F0H        | 80H         | L2_TRANSACTION.S.ANY        | Counts all L2 cache operations.  |         |
| F1H        | 02H         | L2_LINES_IN.S_STATE         | Counts the number of cache lines allocated in the L2 cache in the S (shared) state.  |         |
| F1H        | 04H         | L2_LINES_IN.E_STATE         | Counts the number of cache lines allocated in the L2 cache in the E (exclusive) state.   |         |
| F1H        | 07H         | L2_LINES_IN.ANY             | Counts the number of cache lines allocated in the L2 cache.  |         |
| F2H        | 01H         | L2_LINES_OUT.DEMAND_CLEAN   | Counts L2 clean cache lines evicted by a demand request.   |         |
| F2H        | 02H         | L2_LINES_OUT.DEMAND_DIRTY   | Counts L2 dirty (modified) cache lines evicted by a demand request.  |         |
| F2H        | 04H         | L2_LINES_OUT.PREFETCH_CLEAN | Counts L2 clean cache line evicted by a prefetch request.  |         |
| F2H        | 08H         | L2_LINES_OUT.PREFETCH_DIRTY | Counts L2 modified cache line evicted by a prefetch request.   |         |
| F2H        | 0FH         | L2_LINES_OUT.ANY            | Counts all L2 cache lines evicted for any reason.  |         |
| F4H        | 04H         | SQ_MISC.LRU_HINTS           | Counts number of Super Queue LRU hints sent to L3.   |         |
| F4H        | 10H         | SQ_MISC.SPLIT_LOCK          | Counts the number of SQ lock splits across a cache line.   |         |
| F6H        | 01H         | SQ_FULL_STALL_CYCLES        | Counts cycles the Super Queue is full. Neither of the threads on this core will be able to access the uncore.  |         |
| F7H        | 01H         | FP_ASSIST.ALL               | Counts the number of floating point operations executed that required micro-code assist intervention. Assists are required in the following cases: SSE instructions, (Denormal input when the DAZ flag is off or Underflow result when the FTZ flag is off); x87 instructions, (NaN or denormal are loaded to a register or used as input from memory, Division by 0 or Underflow output). |         |
| F7H        | 02H         | FP_ASSIST.OUTPUT            | Counts number of floating point micro-code assist when the output value (destination register) is invalid.   |         |
| F7H        | 04H         | FP_ASSIST.INPUT             | Counts number of floating point micro-code assist when the input value (one of the source operands to an FP instruction) is invalid.   |         |

**Table 19-21. Performance Events In the Processor Core for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic        | Description   | Comment |
|------------|-------------|----------------------------|---|---------|
| FDH        | 01H         | SIMD_INT_64.PACKED_MPY     | Counts number of SID integer 64 bit packed multiply operations. |         |
| FDH        | 02H         | SIMD_INT_64.PACKED_SHIFT   | Counts number of SID integer 64 bit packed shift operations.    |         |
| FDH        | 04H         | SIMD_INT_64.PACK           | Counts number of SID integer 64 bit pack operations.            |         |
| FDH        | 08H         | SIMD_INT_64.UNPACK         | Counts number of SID integer 64 bit unpack operations.          |         |
| FDH        | 10H         | SIMD_INT_64.PACKED_LOGICAL | Counts number of SID integer 64 bit logical operations.         |         |
| FDH        | 20H         | SIMD_INT_64.PACKED_ARITH   | Counts number of SID integer 64 bit arithmetic operations.      |         |
| FDH        | 40H         | SIMD_INT_64.SHUFFLE_MOVE   | Counts number of SID integer 64 bit shift or move operations.   |         |

Model-specific performance monitoring events of the uncore sub-system for processors with CPUID signature of DisplayFamily\_DisplayModel 06\_25H, 06\_2CH, and 06\_1FH support performance events listed in Table 19-22.

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere**

| Event Num. | Umask Value | Event Mask Mnemonic                        | Description  | Comment |
|------------|-------------|--|--|---------|
| 00H        | 01H         | UNC_GQ_CYCLES_FULL.READ_TRACKER            | Uncore cycles Global Queue read tracker is full.   |         |
| 00H        | 02H         | UNC_GQ_CYCLES_FULL.WRITE_TRACKER           | Uncore cycles Global Queue write tracker is full.  |         |
| 00H        | 04H         | UNC_GQ_CYCLES_FULL.PEER_PROBE_TRACKER      | Uncore cycles Global Queue peer probe tracker is full. The peer probe tracker queue tracks snoops from the IOH and remote sockets.   |         |
| 01H        | 01H         | UNC_GQ_CYCLES_NOT_EMPTY.READ_TRACKER       | Uncore cycles were Global Queue read tracker has at least one valid entry.   |         |
| 01H        | 02H         | UNC_GQ_CYCLES_NOT_EMPTY.WRITE_TRACKER      | Uncore cycles were Global Queue write tracker has at least one valid entry.  |         |
| 01H        | 04H         | UNC_GQ_CYCLES_NOT_EMPTY.PEER_PROBE_TRACKER | Uncore cycles were Global Queue peer probe tracker has at least one valid entry. The peer probe tracker queue tracks IOH and remote socket snoops.   |         |
| 02H        | 01H         | UNC_GQ_OCCUPANCY.READ_TRACKER              | Increments the number of queue entries (code read, data read, and RFOs) in the tread tracker. The GQ read tracker allocate to deallocate occupancy count is divided by the count to obtain the average read tracker latency. |         |
| 03H        | 01H         | UNC_GQ_ALLOC.READ_TRACKER                  | Counts the number of tread tracker allocate to deallocate entries. The GQ read tracker allocate to deallocate occupancy count is divided by the count to obtain the average read tracker latency.                            |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic               | Description  | Comment |
|------------|-------------|-----------------------------------|--|---------|
| 03H        | 02H         | UNC_GQ_ALLOC.RT_L3_MISS           | Counts the number GQ read tracker entries for which a full cache line read has missed the L3. The GQ read tracker L3 miss to fill occupancy count is divided by this count to obtain the average cache line read L3 miss latency. The latency represents the time after which the L3 has determined that the cache line has missed. The time between a GQ read tracker allocation and the L3 determining that the cache line has missed is the average L3 hit latency. The total L3 cache line read miss latency is the hit latency + L3 miss latency. |         |
| 03H        | 04H         | UNC_GQ_ALLOC.RT_TO_L3_RE SP       | Counts the number of GQ read tracker entries that are allocated in the read tracker queue that hit or miss the L3. The GQ read tracker L3 hit occupancy count is divided by this count to obtain the average L3 hit latency.   |         |
| 03H        | 08H         | UNC_GQ_ALLOC.RT_TO_RTID_ ACQUIRED | Counts the number of GQ read tracker entries that are allocated in the read tracker, have missed in the L3 and have not acquired a Request Transaction ID. The GQ read tracker L3 miss to RTID acquired occupancy count is divided by this count to obtain the average latency for a read L3 miss to acquire an RTID.  |         |
| 03H        | 10H         | UNC_GQ_ALLOC.WT_TO_RTID_ ACQUIRED | Counts the number of GQ write tracker entries that are allocated in the write tracker, have missed in the L3 and have not acquired a Request Transaction ID. The GQ write tracker L3 miss to RTID occupancy count is divided by this count to obtain the average latency for a write L3 miss to acquire an RTID.   |         |
| 03H        | 20H         | UNC_GQ_ALLOC.WRITE_TRAC KER       | Counts the number of GQ write tracker entries that are allocated in the write tracker queue that miss the L3. The GQ write tracker occupancy count is divided by this count to obtain the average L3 write miss latency.   |         |
| 03H        | 40H         | UNC_GQ_ALLOC.PEER_PROBE _TRACKER  | Counts the number of GQ peer probe tracker (snoop) entries that are allocated in the peer probe tracker queue that miss the L3. The GQ peer probe occupancy count is divided by this count to obtain the average L3 peer probe miss latency.   |         |
| 04H        | 01H         | UNC_GQ_DATA.FROM_QPI              | Cycles Global Queue Quickpath Interface input data port is busy importing data from the Quickpath Interface. Each cycle the input port can transfer 8 or 16 bytes of data.   |         |
| 04H        | 02H         | UNC_GQ_DATA.FROM_QMC              | Cycles Global Queue Quickpath Memory Interface input data port is busy importing data from the Quickpath Memory Interface. Each cycle the input port can transfer 8 or 16 bytes of data.   |         |
| 04H        | 04H         | UNC_GQ_DATA.FROM_L3               | Cycles GQ L3 input data port is busy importing data from the Last Level Cache. Each cycle the input port can transfer 32 bytes of data.  |         |
| 04H        | 08H         | UNC_GQ_DATA.FROM_CORES_ 02        | Cycles GQ Core 0 and 2 input data port is busy importing data from processor cores 0 and 2. Each cycle the input port can transfer 32 bytes of data.   |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                     | Description   | Comment |
|------------|-------------|---|---|---------|
| 04H        | 10H         | UNC_GQ_DATA.FROM_CORES_13               | Cycles GQ Core 1 and 3 input data port is busy importing data from processor cores 1 and 3. Each cycle the input port can transfer 32 bytes of data.  |         |
| 05H        | 01H         | UNC_GQ_DATA.TO_QPI_QMC                  | Cycles GQ QPI and QMC output data port is busy sending data to the Quickpath Interface or Quickpath Memory Interface. Each cycle the output port can transfer 32 bytes of data.   |         |
| 05H        | 02H         | UNC_GQ_DATA.TO_L3                       | Cycles GQ L3 output data port is busy sending data to the Last Level Cache. Each cycle the output port can transfer 32 bytes of data.   |         |
| 05H        | 04H         | UNC_GQ_DATA.TO_CORES                    | Cycles GQ Core output data port is busy sending data to the Cores. Each cycle the output port can transfer 32 bytes of data.  |         |
| 06H        | 01H         | UNC_SNP_RESP_TO_LOCAL_HOME.I_STATE      | Number of snoop responses to the local home that L3 does not have the referenced cache line.  |         |
| 06H        | 02H         | UNC_SNP_RESP_TO_LOCAL_HOME.S_STATE      | Number of snoop responses to the local home that L3 has the referenced line cached in the S state.  |         |
| 06H        | 04H         | UNC_SNP_RESP_TO_LOCAL_HOME.FWD_S_STATE  | Number of responses to code or data read snoops to the local home that the L3 has the referenced cache line in the E state. The L3 cache line state is changed to the S state and the line is forwarded to the local home in the S state. |         |
| 06H        | 08H         | UNC_SNP_RESP_TO_LOCAL_HOME.FWD_I_STATE  | Number of responses to read invalidate snoops to the local home that the L3 has the referenced cache line in the M state. The L3 cache line state is invalidated and the line is forwarded to the local home in the M state.              |         |
| 06H        | 10H         | UNC_SNP_RESP_TO_LOCAL_HOME.CONFLICT     | Number of conflict snoop responses sent to the local home.  |         |
| 06H        | 20H         | UNC_SNP_RESP_TO_LOCAL_HOME.WB           | Number of responses to code or data read snoops to the local home that the L3 has the referenced line cached in the M state.  |         |
| 07H        | 01H         | UNC_SNP_RESP_TO_REMOTE_HOME.I_STATE     | Number of snoop responses to a remote home that L3 does not have the referenced cache line.   |         |
| 07H        | 02H         | UNC_SNP_RESP_TO_REMOTE_HOME.S_STATE     | Number of snoop responses to a remote home that L3 has the referenced line cached in the S state.   |         |
| 07H        | 04H         | UNC_SNP_RESP_TO_REMOTE_HOME.FWD_S_STATE | Number of responses to code or data read snoops to a remote home that the L3 has the referenced cache line in the E state. The L3 cache line state is changed to the S state and the line is forwarded to the remote home in the S state. |         |
| 07H        | 08H         | UNC_SNP_RESP_TO_REMOTE_HOME.FWD_I_STATE | Number of responses to read invalidate snoops to a remote home that the L3 has the referenced cache line in the M state. The L3 cache line state is invalidated and the line is forwarded to the remote home in the M state.              |         |
| 07H        | 10H         | UNC_SNP_RESP_TO_REMOTE_HOME.CONFLICT    | Number of conflict snoop responses sent to the local home.  |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic              | Description   | Comment |
|------------|-------------|----------------------------------|---|---------|
| 07H        | 20H         | UNC_SNP_RESP_TO_REMOTE_HOME.WB   | Number of responses to code or data read snoops to a remote home that the L3 has the referenced line cached in the M state.   |         |
| 07H        | 24H         | UNC_SNP_RESP_TO_REMOTE_HOME.HITM | Number of HITM snoop responses to a remote home.  |         |
| 08H        | 01H         | UNC_L3_HITS.READ                 | Number of code read, data read and RFO requests that hit in the L3.   |         |
| 08H        | 02H         | UNC_L3_HITS.WRITE                | Number of writeback requests that hit in the L3. Writebacks from the cores will always result in L3 hits due to the inclusive property of the L3.   |         |
| 08H        | 04H         | UNC_L3_HITS.PROBE                | Number of snoops from IOH or remote sockets that hit in the L3.   |         |
| 08H        | 03H         | UNC_L3_HITS.ANY                  | Number of reads and writes that hit the L3.   |         |
| 09H        | 01H         | UNC_L3_MISS.READ                 | Number of code read, data read and RFO requests that miss the L3.   |         |
| 09H        | 02H         | UNC_L3_MISS.WRITE                | Number of writeback requests that miss the L3. Should always be zero as writebacks from the cores will always result in L3 hits due to the inclusive property of the L3.                                      |         |
| 09H        | 04H         | UNC_L3_MISS.PROBE                | Number of snoops from IOH or remote sockets that miss the L3.   |         |
| 09H        | 03H         | UNC_L3_MISS.ANY                  | Number of reads and writes that miss the L3.  |         |
| 0AH        | 01H         | UNC_L3_LINES_IN.M_STATE          | Counts the number of L3 lines allocated in M state. The only time a cache line is allocated in the M state is when the line was forwarded in M state is forwarded due to a Snoop Read Invalidate Own request. |         |
| 0AH        | 02H         | UNC_L3_LINES_IN.E_STATE          | Counts the number of L3 lines allocated in E state.   |         |
| 0AH        | 04H         | UNC_L3_LINES_IN.S_STATE          | Counts the number of L3 lines allocated in S state.   |         |
| 0AH        | 08H         | UNC_L3_LINES_IN.F_STATE          | Counts the number of L3 lines allocated in F state.   |         |
| 0AH        | 0FH         | UNC_L3_LINES_IN.ANY              | Counts the number of L3 lines allocated in any state.   |         |
| 0BH        | 01H         | UNC_L3_LINES_OUT.M_STATE         | Counts the number of L3 lines victimized that were in the M state. When the victim cache line is in M state, the line is written to its home cache agent which can be either local or remote.                 |         |
| 0BH        | 02H         | UNC_L3_LINES_OUT.E_STATE         | Counts the number of L3 lines victimized that were in the E state.  |         |
| 0BH        | 04H         | UNC_L3_LINES_OUT.S_STATE         | Counts the number of L3 lines victimized that were in the S state.  |         |
| 0BH        | 08H         | UNC_L3_LINES_OUT.I_STATE         | Counts the number of L3 lines victimized that were in the I state.  |         |
| 0BH        | 10H         | UNC_L3_LINES_OUT.F_STATE         | Counts the number of L3 lines victimized that were in the F state.  |         |
| 0BH        | 1FH         | UNC_L3_LINES_OUT.ANY             | Counts the number of L3 lines victimized in any state.  |         |
| 0CH        | 01H         | UNC_GQ_SNOOP.GOTO_S              | Counts the number of remote snoops that have requested a cache line be set to the S state.  |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment                                   |
|------------|-------------|--------------------------------|---|---|
| 0CH        | 02H         | UNC_GQ_SNOOP.GOTO_I            | Counts the number of remote snoops that have requested a cache line be set to the I state.                        |   |
| 0CH        | 04H         | UNC_GQ_SNOOP.GOTO_S_HIT_E      | Counts the number of remote snoops that have requested a cache line be set to the S state from E state.           | Requires writing MSR 301H with mask = 2H. |
| 0CH        | 04H         | UNC_GQ_SNOOP.GOTO_S_HIT_F      | Counts the number of remote snoops that have requested a cache line be set to the S state from F (forward) state. | Requires writing MSR 301H with mask = 8H. |
| 0CH        | 04H         | UNC_GQ_SNOOP.GOTO_S_HIT_M      | Counts the number of remote snoops that have requested a cache line be set to the S state from M state.           | Requires writing MSR 301H with mask = 1H. |
| 0CH        | 04H         | UNC_GQ_SNOOP.GOTO_S_HIT_S      | Counts the number of remote snoops that have requested a cache line be set to the S state from S state.           | Requires writing MSR 301H with mask = 4H. |
| 0CH        | 08H         | UNC_GQ_SNOOP.GOTO_I_HIT_E      | Counts the number of remote snoops that have requested a cache line be set to the I state from E state.           | Requires writing MSR 301H with mask = 2H. |
| 0CH        | 08H         | UNC_GQ_SNOOP.GOTO_I_HIT_F      | Counts the number of remote snoops that have requested a cache line be set to the I state from F (forward) state. | Requires writing MSR 301H with mask = 8H. |
| 0CH        | 08H         | UNC_GQ_SNOOP.GOTO_I_HIT_M      | Counts the number of remote snoops that have requested a cache line be set to the I state from M state.           | Requires writing MSR 301H with mask = 1H. |
| 0CH        | 08H         | UNC_GQ_SNOOP.GOTO_I_HIT_S      | Counts the number of remote snoops that have requested a cache line be set to the I state from S state.           | Requires writing MSR 301H with mask = 4H. |
| 20H        | 01H         | UNC_QHL_REQUESTS.IOH_READS     | Counts number of Quickpath Home Logic read requests from the IOH.   |   |
| 20H        | 02H         | UNC_QHL_REQUESTS.IOH_WRITES    | Counts number of Quickpath Home Logic write requests from the IOH.  |   |
| 20H        | 04H         | UNC_QHL_REQUESTS.REMOTE_READS  | Counts number of Quickpath Home Logic read requests from a remote socket.   |   |
| 20H        | 08H         | UNC_QHL_REQUESTS.REMOTE_WRITES | Counts number of Quickpath Home Logic write requests from a remote socket.  |   |
| 20H        | 10H         | UNC_QHL_REQUESTS.LOCAL_READS   | Counts number of Quickpath Home Logic read requests from the local socket.  |   |
| 20H        | 20H         | UNC_QHL_REQUESTS.LOCAL_WRITES  | Counts number of Quickpath Home Logic write requests from the local socket.                                       |   |
| 21H        | 01H         | UNC_QHL_CYCLES_FULL.IOH        | Counts uclk cycles all entries in the Quickpath Home Logic IOH are full.  |   |
| 21H        | 02H         | UNC_QHL_CYCLES_FULL.REMOTE     | Counts uclk cycles all entries in the Quickpath Home Logic remote tracker are full.                               |   |
| 21H        | 04H         | UNC_QHL_CYCLES_FULL.LOCAL      | Counts uclk cycles all entries in the Quickpath Home Logic local tracker are full.                                |   |
| 22H        | 01H         | UNC_QHL_CYCLES_NOT_EMPTY.IOH   | Counts uclk cycles all entries in the Quickpath Home Logic IOH is busy.   |   |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic             | Description  | Comment |
|------------|-------------|---------------------------------|--|---------|
| 22H        | 02H         | UNC_QHL_CYCLES_NOT_EMPTY.REMOTE | Counts uclk cycles all entries in the Quickpath Home Logic remote tracker is busy.   |         |
| 22H        | 04H         | UNC_QHL_CYCLES_NOT_EMPTY.LOCAL  | Counts uclk cycles all entries in the Quickpath Home Logic local tracker is busy.  |         |
| 23H        | 01H         | UNC_QHL_OCCUPANCY.IOH           | QHL IOH tracker allocate to deallocate read occupancy.   |         |
| 23H        | 02H         | UNC_QHL_OCCUPANCY.REMOTE        | QHL remote tracker allocate to deallocate read occupancy.  |         |
| 23H        | 04H         | UNC_QHL_OCCUPANCY.LOCAL         | QHL local tracker allocate to deallocate read occupancy.   |         |
| 24H        | 02H         | UNC_QHL_ADDRESS_CONFLICTS.2WAY  | Counts number of QHL Active Address Table (AAT) entries that saw a max of 2 conflicts. The AAT is a structure that tracks requests that are in conflict. The requests themselves are in the home tracker entries. The count is reported when an AAT entry deallocates. |         |
| 24H        | 04H         | UNC_QHL_ADDRESS_CONFLICTS.3WAY  | Counts number of QHL Active Address Table (AAT) entries that saw a max of 3 conflicts. The AAT is a structure that tracks requests that are in conflict. The requests themselves are in the home tracker entries. The count is reported when an AAT entry deallocates. |         |
| 25H        | 01H         | UNC_QHL_CONFLICT_CYCLES.IOH     | Counts cycles the Quickpath Home Logic IOH Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.   |         |
| 25H        | 02H         | UNC_QHL_CONFLICT_CYCLES.REMOTE  | Counts cycles the Quickpath Home Logic Remote Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.  |         |
| 25H        | 04H         | UNC_QHL_CONFLICT_CYCLES.LOCAL   | Counts cycles the Quickpath Home Logic Local Tracker contains two or more requests with an address conflict. A max of 3 requests can be in conflict.   |         |
| 26H        | 01H         | UNC_QHL_TO_QMC_BYPASS           | Counts number or requests to the Quickpath Memory Controller that bypass the Quickpath Home Logic. All local accesses can be bypassed. For remote requests, only read requests can be bypassed.  |         |
| 28H        | 01H         | UNC_QMC_ISOC_FULL.READ.CH0      | Counts cycles all the entries in the DRAM channel 0 high priority queue are occupied with isochronous read requests.   |         |
| 28H        | 02H         | UNC_QMC_ISOC_FULL.READ.CH1      | Counts cycles all the entries in the DRAM channel 1 high priority queue are occupied with isochronous read requests.   |         |
| 28H        | 04H         | UNC_QMC_ISOC_FULL.READ.CH2      | Counts cycles all the entries in the DRAM channel 2 high priority queue are occupied with isochronous read requests.   |         |
| 28H        | 08H         | UNC_QMC_ISOC_FULL.WRITE.CH0     | Counts cycles all the entries in the DRAM channel 0 high priority queue are occupied with isochronous write requests.  |         |
| 28H        | 10H         | UNC_QMC_ISOC_FULL.WRITE.CH1     | Counts cycles all the entries in the DRAM channel 1 high priority queue are occupied with isochronous write requests.  |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic          | Description  | Comment |
|------------|-------------|------------------------------|--|---------|
| 28H        | 20H         | UNC_QMC_ISOC_FULLL.WRITE.CH2 | Counts cycles all the entries in the DRAM channel 2 high priority queue are occupied with isochronous write requests.  |         |
| 29H        | 01H         | UNC_QMC_BUSY.READ.CH0        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 0.   |         |
| 29H        | 02H         | UNC_QMC_BUSY.READ.CH1        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 1.   |         |
| 29H        | 04H         | UNC_QMC_BUSY.READ.CH2        | Counts cycles where Quickpath Memory Controller has at least 1 outstanding read request to DRAM channel 2.   |         |
| 29H        | 08H         | UNC_QMC_BUSY.WRITE.CH0       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 0.  |         |
| 29H        | 10H         | UNC_QMC_BUSY.WRITE.CH1       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 1.  |         |
| 29H        | 20H         | UNC_QMC_BUSY.WRITE.CH2       | Counts cycles where Quickpath Memory Controller has at least 1 outstanding write request to DRAM channel 2.  |         |
| 2AH        | 01H         | UNC_QMC_OCCUPANCY.CH0        | IMC channel 0 normal read request occupancy.   |         |
| 2AH        | 02H         | UNC_QMC_OCCUPANCY.CH1        | IMC channel 1 normal read request occupancy.   |         |
| 2AH        | 04H         | UNC_QMC_OCCUPANCY.CH2        | IMC channel 2 normal read request occupancy.   |         |
| 2AH        | 07H         | UNC_QMC_OCCUPANCY.ANY        | Normal read request occupancy for any channel.   |         |
| 2BH        | 01H         | UNC_QMC_ISSOC_OCCUPANCY.CH0  | IMC channel 0 issoc read request occupancy.  |         |
| 2BH        | 02H         | UNC_QMC_ISSOC_OCCUPANCY.CH1  | IMC channel 1 issoc read request occupancy.  |         |
| 2BH        | 04H         | UNC_QMC_ISSOC_OCCUPANCY.CH2  | IMC channel 2 issoc read request occupancy.  |         |
| 2BH        | 07H         | UNC_QMC_ISSOC_READS.ANY      | IMC issoc read request occupancy.  |         |
| 2CH        | 01H         | UNC_QMC_NORMAL_READS.CH0     | Counts the number of Quickpath Memory Controller channel 0 medium and low priority read requests. The QMC channel 0 normal read occupancy divided by this count provides the average QMC channel 0 read latency. |         |
| 2CH        | 02H         | UNC_QMC_NORMAL_READS.CH1     | Counts the number of Quickpath Memory Controller channel 1 medium and low priority read requests. The QMC channel 1 normal read occupancy divided by this count provides the average QMC channel 1 read latency. |         |
| 2CH        | 04H         | UNC_QMC_NORMAL_READS.CH2     | Counts the number of Quickpath Memory Controller channel 2 medium and low priority read requests. The QMC channel 2 normal read occupancy divided by this count provides the average QMC channel 2 read latency. |         |



**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                 | Description   | Comment |
|------------|-------------|-------------------------------------|---|---------|
| 2CH        | 07H         | UNC_QMC_NORMAL_READS.ANY            | Counts the number of Quickpath Memory Controller medium and low priority read requests. The QMC normal read occupancy divided by this count provides the average QMC read latency.  |         |
| 2DH        | 01H         | UNC_QMC_HIGH_PRIORITY_READS.CH0     | Counts the number of Quickpath Memory Controller channel 0 high priority isochronous read requests.   |         |
| 2DH        | 02H         | UNC_QMC_HIGH_PRIORITY_READS.CH1     | Counts the number of Quickpath Memory Controller channel 1 high priority isochronous read requests.   |         |
| 2DH        | 04H         | UNC_QMC_HIGH_PRIORITY_READS.CH2     | Counts the number of Quickpath Memory Controller channel 2 high priority isochronous read requests.   |         |
| 2DH        | 07H         | UNC_QMC_HIGH_PRIORITY_READS.ANY     | Counts the number of Quickpath Memory Controller high priority isochronous read requests.   |         |
| 2EH        | 01H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH0 | Counts the number of Quickpath Memory Controller channel 0 critical priority isochronous read requests.   |         |
| 2EH        | 02H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH1 | Counts the number of Quickpath Memory Controller channel 1 critical priority isochronous read requests.   |         |
| 2EH        | 04H         | UNC_QMC_CRITICAL_PRIORITY_READS.CH2 | Counts the number of Quickpath Memory Controller channel 2 critical priority isochronous read requests.   |         |
| 2EH        | 07H         | UNC_QMC_CRITICAL_PRIORITY_READS.ANY | Counts the number of Quickpath Memory Controller critical priority isochronous read requests.   |         |
| 2FH        | 01H         | UNC_QMC_WRITES.FULL.CH0             | Counts number of full cache line writes to DRAM channel 0.  |         |
| 2FH        | 02H         | UNC_QMC_WRITES.FULL.CH1             | Counts number of full cache line writes to DRAM channel 1.  |         |
| 2FH        | 04H         | UNC_QMC_WRITES.FULL.CH2             | Counts number of full cache line writes to DRAM channel 2.  |         |
| 2FH        | 07H         | UNC_QMC_WRITES.FULL.ANY             | Counts number of full cache line writes to DRAM.  |         |
| 2FH        | 08H         | UNC_QMC_WRITES.PARTIAL.CH0          | Counts number of partial cache line writes to DRAM channel 0.   |         |
| 2FH        | 10H         | UNC_QMC_WRITES.PARTIAL.CH1          | Counts number of partial cache line writes to DRAM channel 1.   |         |
| 2FH        | 20H         | UNC_QMC_WRITES.PARTIAL.CH2          | Counts number of partial cache line writes to DRAM channel 2.   |         |
| 2FH        | 38H         | UNC_QMC_WRITES.PARTIAL.ANY          | Counts number of partial cache line writes to DRAM.   |         |
| 30H        | 01H         | UNC_QMC_CANCEL.CH0                  | Counts number of DRAM channel 0 cancel requests.  |         |
| 30H        | 02H         | UNC_QMC_CANCEL.CH1                  | Counts number of DRAM channel 1 cancel requests.  |         |
| 30H        | 04H         | UNC_QMC_CANCEL.CH2                  | Counts number of DRAM channel 2 cancel requests.  |         |
| 30H        | 07H         | UNC_QMC_CANCEL.ANY                  | Counts number of DRAM cancel requests.  |         |
| 31H        | 01H         | UNC_QMC_PRIORITY_UPDATE.S.CH0       | Counts number of DRAM channel 0 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic           | Description   | Comment |
|------------|-------------|-------------------------------|---|---------|
| 31H        | 02H         | UNC_QMC_PRIORITY_UPDATE.S.CH1 | Counts number of DRAM channel 1 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |
| 31H        | 04H         | UNC_QMC_PRIORITY_UPDATE.S.CH2 | Counts number of DRAM channel 2 priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request. |         |
| 31H        | 07H         | UNC_QMC_PRIORITY_UPDATE.S.ANY | Counts number of DRAM priority updates. A priority update occurs when an ISOC high or critical request is received by the QHL and there is a matching request with normal priority that has already been issued to the QMC. In this instance, the QHL will send a priority update to QMC to expedite the request.           |         |
| 32H        | 01H         | UNC_IMC_RETRY.CHO             | Counts number of IMC DRAM channel 0 retries. DRAM retry only occurs when configured in RAS mode.  |         |
| 32H        | 02H         | UNC_IMC_RETRY.CH1             | Counts number of IMC DRAM channel 1 retries. DRAM retry only occurs when configured in RAS mode.  |         |
| 32H        | 04H         | UNC_IMC_RETRY.CH2             | Counts number of IMC DRAM channel 2 retries. DRAM retry only occurs when configured in RAS mode.  |         |
| 32H        | 07H         | UNC_IMC_RETRY.ANY             | Counts number of IMC DRAM retries from any channel. DRAM retry only occurs when configured in RAS mode.   |         |
| 33H        | 01H         | UNC_QHL_FRC_ACK_CNFLTS.IOH    | Counts number of Force Acknowledge Conflict messages sent by the Quickpath Home Logic to the IOH.   |         |
| 33H        | 02H         | UNC_QHL_FRC_ACK_CNFLTS.REMOTE | Counts number of Force Acknowledge Conflict messages sent by the Quickpath Home Logic to the remote home.   |         |
| 33H        | 04H         | UNC_QHL_FRC_ACK_CNFLTS.LOCAL  | Counts number of Force Acknowledge Conflict messages sent by the Quickpath Home Logic to the local home.  |         |
| 33H        | 07H         | UNC_QHL_FRC_ACK_CNFLTS.ANY    | Counts number of Force Acknowledge Conflict messages sent by the Quickpath Home Logic.  |         |
| 34H        | 01H         | UNC_QHL_SLEEPS.IOH_ORDER      | Counts number of occurrences a request was put to sleep due to IOH ordering (write after read) conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.   |         |
| 34H        | 02H         | UNC_QHL_SLEEPS.REMOTE_ORDER   | Counts number of occurrences a request was put to sleep due to remote socket ordering (write after read) conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.   |         |
| 34H        | 04H         | UNC_QHL_SLEEPS.LOCAL_ORDER    | Counts number of occurrences a request was put to sleep due to local socket ordering (write after read) conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.  |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description  | Comment  |
|------------|-------------|---|--|--|
| 34H        | 08H         | UNC_QHL_SLEEPS.IOH_CONFLICT                 | Counts number of occurrences a request was put to sleep due to IOH address conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.  |  |
| 34H        | 10H         | UNC_QHL_SLEEPS.REMOTE_CONFLICT              | Counts number of occurrences a request was put to sleep due to remote socket address conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.  |  |
| 34H        | 20H         | UNC_QHL_SLEEPS.LOCAL_CONFLICT               | Counts number of occurrences a request was put to sleep due to local socket address conflicts. While in the sleep state, the request is not eligible to be scheduled to the QMC.   |  |
| 35H        | 01H         | UNC_ADDR_OPCODE_MATCH.IOH                   | Counts number of requests from the IOH, address/opcode of request is qualified by mask value written to MSR 396H. The following mask values are supported:<br>0: NONE<br>40000000_00000000H:RSPFWDI<br>40001A00_00000000H:RSPFWDS<br>40001D00_00000000H:RSPIWB             | Match opcode/address by writing MSR 396H with mask supported mask value. |
| 35H        | 02H         | UNC_ADDR_OPCODE_MATCH.REMOTE                | Counts number of requests from the remote socket, address/opcode of request is qualified by mask value written to MSR 396H. The following mask values are supported:<br>0: NONE<br>40000000_00000000H:RSPFWDI<br>40001A00_00000000H:RSPFWDS<br>40001D00_00000000H:RSPIWB   | Match opcode/address by writing MSR 396H with mask supported mask value. |
| 35H        | 04H         | UNC_ADDR_OPCODE_MATCH.LOCAL                 | Counts number of requests from the local socket, address/opcode of request is qualified by mask value written to MSR 396H. The following mask values are supported:<br>0: NONE<br>40000000_00000000H:RSPFWDI<br>40001A00_00000000H:RSPFWDS<br>40001D00_00000000H:RSPIWB    | Match opcode/address by writing MSR 396H with mask supported mask value. |
| 40H        | 01H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.HOME.LINK_0  | Counts cycles the Quickpath outbound link 0 HOME virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.  |  |
| 40H        | 02H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.SNOOP.LINK_0 | Counts cycles the Quickpath outbound link 0 SNOOP virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated. |  |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                         | Description   | Comment |
|------------|-------------|---|---|---------|
| 40H        | 04H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.NDR.LINK_0   | Counts cycles the Quickpath outbound link 0 non-data response virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.    |         |
| 40H        | 08H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.HOME.LINK_1  | Counts cycles the Quickpath outbound link 1 HOME virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                 |         |
| 40H        | 10H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.SNOOP.LINK_1 | Counts cycles the Quickpath outbound link 1 SNOOP virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                |         |
| 40H        | 20H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.NDR.LINK_1   | Counts cycles the Quickpath outbound link 1 non-data response virtual channel is stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.    |         |
| 40H        | 07H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.LINK_0       | Counts cycles the Quickpath outbound link 0 virtual channels are stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                    |         |
| 40H        | 38H         | UNC_QPI_TX_STALLED_SINGLE_FLIT.LINK_1       | Counts cycles the Quickpath outbound link 1 virtual channels are stalled due to lack of a VNA and VNO credit. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                    |         |
| 41H        | 01H         | UNC_QPI_TX_STALLED_MULTIFLIT.DRS.LINK_0     | Counts cycles the Quickpath outbound link 0 Data Response virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.         |         |
| 41H        | 02H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCB.LINK_0     | Counts cycles the Quickpath outbound link 0 Non-Coherent Bypass virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |
| 41H        | 04H         | UNC_QPI_TX_STALLED_MULTIFLIT.NCS.LINK_0     | Counts cycles the Quickpath outbound link 0 Non-Coherent Standard virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated. |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                      | Description   | Comment |
|------------|-------------|--|---|---------|
| 41H        | 08H         | UNC_QPI_TX_STALLED_MULTI_FLIT.DRS.LINK_1 | Counts cycles the Quickpath outbound link 1 Data Response virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.         |         |
| 41H        | 10H         | UNC_QPI_TX_STALLED_MULTI_FLIT.NCB.LINK_1 | Counts cycles the Quickpath outbound link 1 Non-Coherent Bypass virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.   |         |
| 41H        | 20H         | UNC_QPI_TX_STALLED_MULTI_FLIT.NCS.LINK_1 | Counts cycles the Quickpath outbound link 1 Non-Coherent Standard virtual channel is stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated. |         |
| 41H        | 07H         | UNC_QPI_TX_STALLED_MULTI_FLIT.LINK_0     | Counts cycles the Quickpath outbound link 0 virtual channels are stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                     |         |
| 41H        | 38H         | UNC_QPI_TX_STALLED_MULTI_FLIT.LINK_1     | Counts cycles the Quickpath outbound link 1 virtual channels are stalled due to lack of VNA and VNO credits. Note that this event does not filter out when a flit would not have been selected for arbitration because another virtual channel is getting arbitrated.                     |         |
| 42H        | 01H         | UNC_QPI_TX_HEADER.FULL.LINK_0            | Number of cycles that the header buffer in the Quickpath Interface outbound link 0 is full.   |         |
| 42H        | 02H         | UNC_QPI_TX_HEADER.BUSY.LINK_0            | Number of cycles that the header buffer in the Quickpath Interface outbound link 0 is busy.   |         |
| 42H        | 04H         | UNC_QPI_TX_HEADER.FULL.LINK_1            | Number of cycles that the header buffer in the Quickpath Interface outbound link 1 is full.   |         |
| 42H        | 08H         | UNC_QPI_TX_HEADER.BUSY.LINK_1            | Number of cycles that the header buffer in the Quickpath Interface outbound link 1 is busy.   |         |
| 43H        | 01H         | UNC_QPI_RX_NO_PPT_CREDIT.STALLS.LINK_0   | Number of cycles that snoop packets incoming to the Quickpath Interface link 0 are stalled and not sent to the GQ because the GQ Peer Probe Tracker (PPT) does not have any available entries.  |         |
| 43H        | 02H         | UNC_QPI_RX_NO_PPT_CREDIT.STALLS.LINK_1   | Number of cycles that snoop packets incoming to the Quickpath Interface link 1 are stalled and not sent to the GQ because the GQ Peer Probe Tracker (PPT) does not have any available entries.  |         |
| 60H        | 01H         | UNC_DRAM_OPEN.CHO                        | Counts number of DRAM Channel 0 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic            | Description   | Comment |
|------------|-------------|--------------------------------|---|---------|
| 60H        | 02H         | UNC_DRAM_OPEN.CH1              | Counts number of DRAM Channel 1 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |
| 60H        | 04H         | UNC_DRAM_OPEN.CH2              | Counts number of DRAM Channel 2 open commands issued either for read or write. To read or write data, the referenced DRAM page must first be opened.  |         |
| 61H        | 01H         | UNC_DRAM_PAGE_CLOSE.CH0        | DRAM channel 0 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |
| 61H        | 02H         | UNC_DRAM_PAGE_CLOSE.CH1        | DRAM channel 1 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |
| 61H        | 04H         | UNC_DRAM_PAGE_CLOSE.CH2        | DRAM channel 2 command issued to CLOSE a page due to page idle timer expiration. Closing a page is done by issuing a precharge.   |         |
| 62H        | 01H         | UNC_DRAM_PAGE_MISS.CH0         | Counts the number of precharges (PRE) that were issued to DRAM channel 0 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 62H        | 02H         | UNC_DRAM_PAGE_MISS.CH1         | Counts the number of precharges (PRE) that were issued to DRAM channel 1 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 62H        | 04H         | UNC_DRAM_PAGE_MISS.CH2         | Counts the number of precharges (PRE) that were issued to DRAM channel 2 because there was a page miss. A page miss refers to a situation in which a page is currently open and another page from the same bank needs to be opened. The new page experiences a page miss. Closing of the old page is done by issuing a precharge. |         |
| 63H        | 01H         | UNC_DRAM_READ_CAS.CH0          | Counts the number of times a read CAS command was issued on DRAM channel 0.   |         |
| 63H        | 02H         | UNC_DRAM_READ_CAS.AUTO PRE_CH0 | Counts the number of times a read CAS command was issued on DRAM channel 0 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 63H        | 04H         | UNC_DRAM_READ_CAS.CH1          | Counts the number of times a read CAS command was issued on DRAM channel 1.   |         |
| 63H        | 08H         | UNC_DRAM_READ_CAS.AUTO PRE_CH1 | Counts the number of times a read CAS command was issued on DRAM channel 1 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 63H        | 10H         | UNC_DRAM_READ_CAS.CH2          | Counts the number of times a read CAS command was issued on DRAM channel 2.   |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                | Description  | Comment |
|------------|-------------|------------------------------------|--|---------|
| 63H        | 20H         | UNC_DRAM_READ_CAS.AUTO PRE_CH2     | Counts the number of times a read CAS command was issued on DRAM channel 2 where the command issued used the auto-precharge (auto page close) mode.  |         |
| 64H        | 01H         | UNC_DRAM_WRITE_CAS.CHO             | Counts the number of times a write CAS command was issued on DRAM channel 0.   |         |
| 64H        | 02H         | UNC_DRAM_WRITE_CAS.AUTO PRE_CH0    | Counts the number of times a write CAS command was issued on DRAM channel 0 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 64H        | 04H         | UNC_DRAM_WRITE_CAS.CH1             | Counts the number of times a write CAS command was issued on DRAM channel 1.   |         |
| 64H        | 08H         | UNC_DRAM_WRITE_CAS.AUTO PRE_CH1    | Counts the number of times a write CAS command was issued on DRAM channel 1 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 64H        | 10H         | UNC_DRAM_WRITE_CAS.CH2             | Counts the number of times a write CAS command was issued on DRAM channel 2.   |         |
| 64H        | 20H         | UNC_DRAM_WRITE_CAS.AUTO PRE_CH2    | Counts the number of times a write CAS command was issued on DRAM channel 2 where the command issued used the auto-precharge (auto page close) mode.   |         |
| 65H        | 01H         | UNC_DRAM_REFRESH.CHO               | Counts number of DRAM channel 0 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 65H        | 02H         | UNC_DRAM_REFRESH.CH1               | Counts number of DRAM channel 1 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 65H        | 04H         | UNC_DRAM_REFRESH.CH2               | Counts number of DRAM channel 2 refresh commands. DRAM loses data content over time. In order to keep correct data content, the data values have to be refreshed periodically.                   |         |
| 66H        | 01H         | UNC_DRAM_PRE_ALL.CHO               | Counts number of DRAM Channel 0 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |
| 66H        | 02H         | UNC_DRAM_PRE_ALL.CH1               | Counts number of DRAM Channel 1 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |
| 66H        | 04H         | UNC_DRAM_PRE_ALL.CH2               | Counts number of DRAM Channel 2 precharge-all (PREALL) commands that close all open pages in a rank. PREALL is issued when the DRAM needs to be refreshed or needs to go into a power down mode. |         |
| 67H        | 01H         | UNC_DRAM_THERMAL_THROT TLED        | Uncore cycles DRAM was throttled due to its temperature being above the thermal throttling threshold.  |         |
| 80H        | 01H         | UNC_THERMAL_THROTTLING_TEMP.CORE_0 | Cycles that the PCU records that core 0 is above the thermal throttling threshold temperature.   |         |

**Table 19-22. Performance Events In the Processor Uncore for Processors Based on Intel® Microarchitecture Code Name Westmere (Contd.)**

| Event Num. | Umask Value | Event Mask Mnemonic                   | Description   | Comment |
|------------|-------------|---------------------------------------|---|---------|
| 80H        | 02H         | UNC_THERMAL_THROTTLING_TEMP.CORE_1    | Cycles that the PCU records that core 1 is above the thermal throttling threshold temperature.  |         |
| 80H        | 04H         | UNC_THERMAL_THROTTLING_TEMP.CORE_2    | Cycles that the PCU records that core 2 is above the thermal throttling threshold temperature.  |         |
| 80H        | 08H         | UNC_THERMAL_THROTTLING_TEMP.CORE_3    | Cycles that the PCU records that core 3 is above the thermal throttling threshold temperature.  |         |
| 81H        | 01H         | UNC_THERMAL_THROTTLED_TEMP.CORE_0     | Cycles that the PCU records that core 0 is in the power throttled state due to core's temperature being above the thermal throttling threshold.       |         |
| 81H        | 02H         | UNC_THERMAL_THROTTLED_TEMP.CORE_1     | Cycles that the PCU records that core 1 is in the power throttled state due to core's temperature being above the thermal throttling threshold.       |         |
| 81H        | 04H         | UNC_THERMAL_THROTTLED_TEMP.CORE_2     | Cycles that the PCU records that core 2 is in the power throttled state due to core's temperature being above the thermal throttling threshold.       |         |
| 81H        | 08H         | UNC_THERMAL_THROTTLED_TEMP.CORE_3     | Cycles that the PCU records that core 3 is in the power throttled state due to core's temperature being above the thermal throttling threshold.       |         |
| 82H        | 01H         | UNC_PROCHOT_ASSERTION                 | Number of system assertions of PROCHOT indicating the entire processor has exceeded the thermal limit.  |         |
| 83H        | 01H         | UNC_THERMAL_THROTTLING_PROCHOT.CORE_0 | Cycles that the PCU records that core 0 is a low power state due to the system asserting PROCHOT the entire processor has exceeded the thermal limit. |         |
| 83H        | 02H         | UNC_THERMAL_THROTTLING_PROCHOT.CORE_1 | Cycles that the PCU records that core 1 is a low power state due to the system asserting PROCHOT the entire processor has exceeded the thermal limit. |         |
| 83H        | 04H         | UNC_THERMAL_THROTTLING_PROCHOT.CORE_2 | Cycles that the PCU records that core 2 is a low power state due to the system asserting PROCHOT the entire processor has exceeded the thermal limit. |         |
| 83H        | 08H         | UNC_THERMAL_THROTTLING_PROCHOT.CORE_3 | Cycles that the PCU records that core 3 is a low power state due to the system asserting PROCHOT the entire processor has exceeded the thermal limit. |         |
| 84H        | 01H         | UNC_TURBO_MODE.CORE_0                 | Uncore cycles that core 0 is operating in turbo mode.   |         |
| 84H        | 02H         | UNC_TURBO_MODE.CORE_1                 | Uncore cycles that core 1 is operating in turbo mode.   |         |
| 84H        | 04H         | UNC_TURBO_MODE.CORE_2                 | Uncore cycles that core 2 is operating in turbo mode.   |         |
| 84H        | 08H         | UNC_TURBO_MODE.CORE_3                 | Uncore cycles that core 3 is operating in turbo mode.   |         |
| 85H        | 02H         | UNC_CYCLES_UNHALTED_L3_FLL_ENABLE     | Uncore cycles that at least one core is unhalted and all L3 ways are enabled.   |         |
| 86H        | 01H         | UNC_CYCLES_UNHALTED_L3_FLL_DISABLE    | Uncore cycles that at least one core is unhalted and all L3 ways are disabled.  |         |



## 19.11 PERFORMANCE MONITORING EVENTS FOR INTEL® XEON® PROCESSOR 5200, 5400 SERIES AND INTEL® CORE™ 2 EXTREME PROCESSORS QX 9000 SERIES

Processors based on the Enhanced Intel Core microarchitecture support the architectural and model-specific performance monitoring events listed in Table 19-1 and Table 19-25. In addition, they also support the following model-specific performance monitoring events listed in Table 19-23. Fixed counters support the architecture events defined in Table 19-24.

**Table 19-23. Performance Events for Processors Based on Enhanced Intel Core Microarchitecture**

| Event Num. | Umask Value | Event Mask Mnemonic                       | Description  | Comment |
|------------|-------------|---|--|---------|
| COH        | 08H         | INST_RETIRED.VM_HOST                      | Instruction retired while in VMX root operations.  |         |
| D2H        | 10H         | RAT_STAALS.OTHER_SERIALI<br>ZATION_STALLS | This event counts the number of stalls due to other RAT resource serialization not counted by Umask value 0FH. |         |

## 19.12 PERFORMANCE MONITORING EVENTS FOR INTEL® XEON® PROCESSOR 3000, 3200, 5100, 5300 SERIES AND INTEL® CORE™ 2 DUO PROCESSORS

Processors based on the Intel® Core™ microarchitecture support architectural and model-specific performance monitoring events.

Fixed-function performance counters are introduced first on processors based on Intel Core microarchitecture. Table 19-24 lists pre-defined performance events that can be counted using fixed-function performance counters.

**Table 19-24. Fixed-Function Performance Counter and Pre-defined Performance Events**

| Fixed-Function Performance Counter       | Address | Event Mask Mnemonic   | Description   |
|--|---------|-----------------------|---|
| MSR_PERF_FIXED_CTR0/IA32_PERF_FIXED_CTR0 | 309H    | Inst_Retired.Any      | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers.   |
| MSR_PERF_FIXED_CTR1/IA32_PERF_FIXED_CTR1 | 30AH    | CPU_CLK_UNHALTED.CORE | This event counts the number of core cycles while the core is not in a halt state. The core enters the halt state when it is running the HLT instruction. This event is a component in many key event ratios.<br><br>The core frequency may change from time to time due to transitions associated with Enhanced Intel SpeedStep Technology or TM2. For this reason this event may have a changing ratio with regards to time. When the core frequency is constant, this event can approximate elapsed time while the core was not in halt state. |
| MSR_PERF_FIXED_CTR2/IA32_PERF_FIXED_CTR2 | 30BH    | CPU_CLK_UNHALTED.REF  | This event counts the number of reference cycles when the core is not in a halt state and not in a TM stop-clock state. The core enters the halt state when it is running the HLT instruction or the MWAIT instruction.   |

**Table 19-24. Fixed-Function Performance Counter and Pre-defined Performance Events (Contd.)**

| Fixed-Function Performance Counter | Address | Event Mask Mnemonic | Description   |
|------------------------------------|---------|---------------------|---|
|                                    |         |                     | <p>This event is not affected by core frequency changes (e.g., P states) but counts at the same frequency as the time stamp counter. This event can approximate elapsed time while the core was not in halt state and not in a TM stop-clock state.</p> <p>This event has a constant ratio with the CPU_CLK_UNHALTED.BUS event.</p> |

Table 19-25 lists general-purpose model-specific performance monitoring events supported in processors based on Intel® Core™ microarchitecture. For convenience, Table 19-25 also includes architectural events and describes minor model-specific behavior where applicable. Software must use a general-purpose performance counter to count events listed in Table 19-25.

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture**

| Event Num | Umask Value | Event Name               | Definition   | Description and Comment   |
|-----------|-------------|--------------------------|--|---|
| 03H       | 02H         | LOAD_BLOCK.STA           | Loads blocked by a preceding store with unknown address.                                 | <p>This event indicates that loads are blocked by preceding stores. A load is blocked when there is a preceding store to an address that is not yet calculated. The number of events is greater or equal to the number of load operations that were blocked.</p> <p>If the load and the store are always to different addresses, check why the memory disambiguation mechanism is not working. To avoid such blocks, increase the distance between the store and the following load so that the store address is known at the time the load is dispatched.</p>  |
| 03H       | 04H         | LOAD_BLOCK.STD           | Loads blocked by a preceding store with unknown data.                                    | <p>This event indicates that loads are blocked by preceding stores. A load is blocked when there is a preceding store to the same address and the stored data value is not yet known. The number of events is greater or equal to the number of load operations that were blocked.</p> <p>To avoid such blocks, increase the distance between the store and the dependent load, so that the store data is known at the time the load is dispatched.</p>   |
| 03H       | 08H         | LOAD_BLOCK.OVERLAP_STORE | Loads that partially overlap an earlier store, or 4-Kbyte aliased with a previous store. | <p>This event indicates that loads are blocked due to a variety of reasons. Some of the triggers for this event are when a load is blocked by a preceding store, in one of the following:</p> <ul style="list-style-type: none"> <li>▪ Some of the loaded byte locations are written by the preceding store and some are not.</li> <li>▪ The load is from bytes written by the preceding store, the store is aligned to its size and either: <ul style="list-style-type: none"> <li>▪ The load's data size is one or two bytes and it is not aligned to the store.</li> <li>▪ The load's data size is of four or eight bytes and the load is misaligned.</li> </ul> </li> </ul> |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name              | Definition   | Description and Comment  |
|-----------|-------------|-------------------------|--|--|
|           |             |                         |  | <ul style="list-style-type: none"> <li>The load is from bytes written by the preceding store, the store is misaligned and the load is not aligned on the beginning of the store.</li> <li>The load is split over an eight byte boundary (excluding 16-byte loads).</li> <li>The load and store have the same offset relative to the beginning of different 4-KByte pages. This case is also called 4-KByte aliasing.</li> <li>In all these cases the load is blocked until after the blocking store retires and the stored data is committed to the cache hierarchy.</li> </ul>  |
| 03H       | 10H         | LOAD_BLOCK.UNTIL_RETIRE | Loads blocked until retirement.  | This event indicates that load operations were blocked until retirement. The number of events is greater or equal to the number of load operations that were blocked. This includes mainly uncacheable loads and split loads (loads that cross the cache line boundary) but may include other cases where loads are blocked until retirement.  |
| 03H       | 20H         | LOAD_BLOCK.L1D          | Loads blocked by the L1 data cache.  | This event indicates that loads are blocked due to one or more reasons. Some triggers for this event are: <ul style="list-style-type: none"> <li>The number of L1 data cache misses exceeds the maximum number of outstanding misses supported by the processor. This includes misses generated as result of demand fetches, software prefetches or hardware prefetches.</li> <li>Cache line split loads.</li> <li>Partial reads, such as reads to un-cacheable memory, I/O instructions and more.</li> <li>A locked load operation is in progress. The number of events is greater or equal to the number of load operations that were blocked.</li> </ul>      |
| 04H       | 01H         | SB_DRAIN_CYCLES         | Cycles while stores are blocked due to store buffer drain.                   | This event counts every cycle during which the store buffer is draining. This includes: <ul style="list-style-type: none"> <li>Serializing operations such as CPUID</li> <li>Synchronizing operations such as XCHG</li> <li>Interrupt acknowledgment</li> <li>Other conditions, such as cache flushing</li> </ul>  |
| 04H       | 02H         | STORE_BLOCK.ORDER       | Cycles while store is waiting for a preceding store to be globally observed. | This event counts the total duration, in number of cycles, which stores are waiting for a preceding stored cache line to be observed by other cores. This situation happens as a result of the strong store ordering behavior, as defined in "Memory Ordering," Chapter 8, <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A</i> .<br><br>The stall may occur and be noticeable if there are many cases when a store either misses the L1 data cache or hits a cache line in the Shared state. If the store requires a bus transaction to read the cache line then the stall ends when snoop response for the bus transaction arrives. |
| 04H       | 08H         | STORE_BLOCK.SNOOP       | A store is blocked due to a conflict with an external or internal snoop.     | This event counts the number of cycles the store port was used for snooping the L1 data cache and a store was stalled by the snoop. The store is typically resubmitted one cycle later.  |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name             | Definition   | Description and Comment   |
|-----------|-------------|------------------------|--|---|
| 06H       | 00H         | SEGMENT_REG_LOADS      | Number of segment register loads.  | <p>This event counts the number of segment register load operations. Instructions that load new values into segment registers cause a penalty.</p> <p>This event indicates performance issues in 16-bit code. If this event occurs frequently, it may be useful to calculate the number of instructions retired per segment register load. If the resulting calculation is low (on average a small number of instructions are executed between segment register loads), then the code's segment register usage should be optimized.</p> <p>As a result of branch misprediction, this event is speculative and may include segment register loads that do not actually occur. However, most segment register loads are internally serialized and such speculative effects are minimized.</p> |
| 07H       | 00H         | SSE_PRE_EXEC.NTA       | Streaming SIMD Extensions (SSE) Prefetch NTA instructions executed.              | <p>This event counts the number of times the SSE instruction prefetchNTA is executed.</p> <p>This instruction prefetches the data to the L1 data cache.</p>   |
| 07H       | 01H         | SSE_PRE_EXECL1         | Streaming SIMD Extensions (SSE) PrefetchT0 instructions executed.                | This event counts the number of times the SSE instruction prefetchT0 is executed. This instruction prefetches the data to the L1 data cache and L2 cache.   |
| 07H       | 02H         | SSE_PRE_EXECL2         | Streaming SIMD Extensions (SSE) PrefetchT1 and PrefetchT2 instructions executed. | This event counts the number of times the SSE instructions prefetchT1 and prefetchT2 are executed. These instructions prefetch the data to the L2 cache.  |
| 07H       | 03H         | SSE_PRE_EXEC.STORES    | Streaming SIMD Extensions (SSE) Weakly-ordered store instructions executed.      | This event counts the number of times SSE non-temporal store instructions are executed.   |
| 08H       | 01H         | DTLB_MISSES.ANY        | Memory accesses that missed the DTLB.  | <p>This event counts the number of Data Table Lookaside Buffer (DTLB) misses. The count includes misses detected as a result of speculative accesses.</p> <p>Typically a high count for this event indicates that the code accesses a large number of data pages.</p>   |
| 08H       | 02H         | DTLB_MISSES.MISS_LD    | DTLB misses due to load operations.  | <p>This event counts the number of Data Table Lookaside Buffer (DTLB) misses due to load operations.</p> <p>This count includes misses detected as a result of speculative accesses.</p>  |
| 08H       | 04H         | DTLB_MISSES.LO_MISS_LD | LO DTLB misses due to load operations.   | <p>This event counts the number of level 0 Data Table Lookaside Buffer (DTLB0) misses due to load operations.</p> <p>This count includes misses detected as a result of speculative accesses. Loads that miss that DTLB0 and hit the DTLB1 can incur two-cycle penalty.</p>   |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                    | Definition                                       | Description and Comment  |
|-----------|-------------|-------------------------------|--|--|
| 08H       | 08H         | DTLB_MISSES.MISS_ST           | TLB misses due to store operations.              | This event counts the number of Data Table Lookaside Buffer (DTLB) misses due to store operations.<br><br>This count includes misses detected as a result of speculative accesses. Address translation for store operations is performed in the DTLB1.   |
| 09H       | 01H         | MEMORY_DISAMBIGUATION.RESET   | Memory disambiguation reset cycles.              | This event counts the number of cycles during which memory disambiguation misprediction occurs. As a result the execution pipeline is cleaned and execution of the mispredicted load instruction and all succeeding instructions restarts.<br><br>This event occurs when the data address accessed by a load instruction, collides infrequently with preceding stores, but usually there is no collision. It happens rarely, and may have a penalty of about 20 cycles.  |
| 09H       | 02H         | MEMORY_DISAMBIGUATION.SUCCESS | Number of loads successfully disambiguated.      | This event counts the number of load operations that were successfully disambiguated. Loads are preceded by a store with an unknown address, but they are not blocked.   |
| 0CH       | 01H         | PAGE_WALKS.COUNT              | Number of page-walks executed.                   | This event counts the number of page-walks executed due to either a DTLB or ITLB miss.<br><br>The page walk duration, PAGE_WALKS.CYCLES, divided by number of page walks is the average duration of a page walk. The average can hint whether most of the page-walks are satisfied by the caches or cause an L2 cache miss.  |
| 0CH       | 02H         | PAGE_WALKS.CYCLES             | Duration of page-walks in core cycles.           | This event counts the duration of page-walks in core cycles. The paging mode in use typically affects the duration of page walks.<br><br>Page walk duration divided by number of page walks is the average duration of page-walks. The average can hint at whether most of the page-walks are satisfied by the caches or cause an L2 cache miss.   |
| 10H       | 00H         | FP_COMP_OPS_EXE               | Floating point computational micro-ops executed. | This event counts the number of floating point computational micro-ops executed.<br>Use IA32_PMC0 only.  |
| 11H       | 00H         | FP_ASSIST                     | Floating point assists.                          | This event counts the number of floating point operations executed that required micro-code assist intervention. Assists are required in the following cases: <ul style="list-style-type: none"> <li>▪ Streaming SIMD Extensions (SSE) instructions:</li> <li>▪ Denormal input when the DAZ (Denormals Are Zeros) flag is off</li> <li>▪ Underflow result when the FTZ (Flush To Zero) flag is off</li> <li>▪ X87 instructions:</li> <li>▪ NaN or denormal are loaded to a register or used as input from memory</li> <li>▪ Division by 0</li> <li>▪ Underflow output</li> </ul> Use IA32_PMC1 only. |
| 12H       | 00H         | MUL                           | Multiply operations executed.                    | This event counts the number of multiply operations executed. This includes integer as well as floating point multiply operations.<br>Use IA32_PMC1 only.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value     | Event Name             | Definition   | Description and Comment  |
|-----------|-----------------|------------------------|--|--|
| 13H       | 00H             | DIV                    | Divide operations executed.  | This event counts the number of divide operations executed. This includes integer divides, floating point divides and square-root operations executed.<br>Use IA32_PMC1 only.  |
| 14H       | 00H             | CYCLES_DIV_BUSY        | Cycles the divider busy.   | This event counts the number of cycles the divider is busy executing divide or square root operations. The divide can be integer, X87 or Streaming SIMD Extensions (SSE). The square root operation can be either X87 or SSE.<br>Use IA32_PMC0 only.   |
| 18H       | 00H             | IDLE_DURING_DIV        | Cycles the divider is busy and all other execution units are idle. | This event counts the number of cycles the divider is busy (with a divide or a square root operation) and no other execution unit or load operation is in progress.<br>Load operations are assumed to hit the L1 data cache. This event considers only micro-ops dispatched after the divider started operating.<br>Use IA32_PMC0 only.                      |
| 19H       | 00H             | DELAYED_BYPASS.FP      | Delayed bypass to FP operation.                                    | This event counts the number of times floating point operations use data immediately after the data was generated by a non-floating point execution unit. Such cases result in one penalty cycle due to data bypass between the units.<br>Use IA32_PMC1 only.  |
| 19H       | 01H             | DELAYED_BYPASS.SIMD    | Delayed bypass to SIMD operation.                                  | This event counts the number of times SIMD operations use data immediately after the data was generated by a non-SIMD execution unit. Such cases result in one penalty cycle due to data bypass between the units.<br>Use IA32_PMC1 only.  |
| 19H       | 02H             | DELAYED_BYPASS.LOAD    | Delayed bypass to load operation.                                  | This event counts the number of delayed bypass penalty cycles that a load operation incurred.<br>When load operations use data immediately after the data was generated by an integer execution unit, they may (pending on certain dynamic internal conditions) incur one penalty cycle due to delayed data bypass between the units.<br>Use IA32_PMC1 only. |
| 21H       | See Table 18-61 | L2_ADS.(Core)          | Cycles L2 address bus is in use.                                   | This event counts the number of cycles the L2 address bus is being used for accesses to the L2 cache or bus queue. It can count occurrences for this core or both cores.   |
| 23H       | See Table 18-61 | L2_DBUS_BUSY_RD.(Core) | Cycles the L2 transfers data to the core.                          | This event counts the number of cycles during which the L2 data bus is busy transferring data from the L2 cache to the core. It counts for all L1 cache misses (data and instruction) that hit the L2 cache.<br>This event can count occurrences for this core or both cores.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value  | Event Name                               | Definition                                | Description and Comment  |
|-----------|--|--|---|--|
| 24H       | Combined mask from Table 18-61 and Table 18-63               | L2_LINES_IN.<br>(Core, Prefetch)         | L2 cache misses.                          | This event counts the number of cache lines allocated in the L2 cache. Cache lines are allocated in the L2 cache as a result of requests from the L1 data and instruction caches and the L2 hardware prefetchers to cache lines that are missing in the L2 cache.<br><br>This event can count occurrences for this core or both cores. It can also count demand requests and L2 hardware prefetch requests together or separately. |
| 25H       | See Table 18-61  | L2_M_LINES_IN.<br>(Core)                 | L2 cache line modifications.              | This event counts whenever a modified cache line is written back from the L1 data cache to the L2 cache.<br><br>This event can count occurrences for this core or both cores.  |
| 26H       | See Table 18-61 and Table 18-63                              | L2_LINES_OUT.<br>(Core, Prefetch)        | L2 cache lines evicted.                   | This event counts the number of L2 cache lines evicted.<br><br>This event can count occurrences for this core or both cores. It can also count evictions due to demand requests and L2 hardware prefetch requests together or separately.  |
| 27H       | See Table 18-61 and Table 18-63                              | L2_M_LINES_OUT.(Core, Prefetch)          | Modified lines evicted from the L2 cache. | This event counts the number of L2 modified cache lines evicted. These lines are written back to memory unless they also exist in a modified-state in one of the L1 data caches.<br><br>This event can count occurrences for this core or both cores. It can also count evictions due to demand requests and L2 hardware prefetch requests together or separately.   |
| 28H       | Combined mask from Table 18-61 and Table 18-64               | L2_IFETCH.(Core, Cache Line State)       | L2 cacheable instruction fetch requests.  | This event counts the number of instruction cache line requests from the IFU. It does not include fetch requests from uncacheable memory. It does not include ITLB miss accesses.<br><br>This event can count occurrences for this core or both cores. It can also count accesses to cache lines at different MESI states.   |
| 29H       | Combined mask from Table 18-61, Table 18-63, and Table 18-64 | L2_LD.(Core, Prefetch, Cache Line State) | L2 cache reads.                           | This event counts L2 cache read requests coming from the L1 data cache and L2 prefetchers.<br><br>The event can count occurrences: <ul style="list-style-type: none"> <li>▪ For this core or both cores.</li> <li>▪ Due to demand requests and L2 hardware prefetch requests together or separately.</li> <li>▪ Of accesses to cache lines at different MESI states.</li> </ul>  |
| 2AH       | See Table 18-61 and Table 18-64                              | L2_ST.(Core, Cache Line State)           | L2 store requests.                        | This event counts all store operations that miss the L1 data cache and request the data from the L2 cache.<br><br>The event can count occurrences for this core or both cores. It can also count accesses to cache lines at different MESI states.   |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value                                   | Event Name  | Definition  | Description and Comment   |
|-----------|---|---|---|---|
| 2BH       | See Table 18-61 and Table 18-64               | L2_LOCK.(Core, Cache Line State)                  | L2 locked accesses.   | This event counts all locked accesses to cache lines that miss the L1 data cache.<br><br>The event can count occurrences for this core or both cores. It can also count accesses to cache lines at different MESI states.   |
| 2EH       | See Table 18-61, Table 18-63, and Table 18-64 | L2_RQSTS.(Core, Prefetch, Cache Line State)       | L2 cache requests.  | This event counts all completed L2 cache requests. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, instruction fetches, and all L2 hardware prefetch requests.<br><br>This event can count occurrences: <ul style="list-style-type: none"> <li>For this core or both cores.</li> <li>Due to demand requests and L2 hardware prefetch requests together, or separately.</li> <li>Of accesses to cache lines at different MESI states.</li> </ul>  |
| 2EH       | 41H   | L2_RQSTS.SELF.DEMAND.I_STATE                      | L2 cache demand requests from this core that missed the L2.       | This event counts all completed L2 cache demand requests from this core that miss the L2 cache. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, and instruction fetches.<br><br>This is an architectural performance event.  |
| 2EH       | 4FH   | L2_RQSTS.SELF.DEMAND.MESI                         | L2 cache demand requests from this core.                          | This event counts all completed L2 cache demand requests from this core. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, and instruction fetches.<br><br>This is an architectural performance event.   |
| 30H       | See Table 18-61, Table 18-63, and Table 18-64 | L2_REJECT_BUSQ.(Core, Prefetch, Cache Line State) | Rejected L2 cache requests.                                       | This event indicates that a pending L2 cache request that requires a bus transaction is delayed from moving to the bus queue. Some of the reasons for this event are: <ul style="list-style-type: none"> <li>The bus queue is full.</li> <li>The bus queue already holds an entry for a cache line in the same set.</li> </ul> The number of events is greater or equal to the number of requests that were rejected. <ul style="list-style-type: none"> <li>For this core or both cores.</li> <li>Due to demand requests and L2 hardware prefetch requests together, or separately.</li> <li>Of accesses to cache lines at different MESI states.</li> </ul> |
| 32H       | See Table 18-61                               | L2_NO_REQ.(Core)                                  | Cycles no L2 cache requests are pending.                          | This event counts the number of cycles that no L2 cache requests were pending from a core. When using the BOTH_CORE modifier, the event counts only if none of the cores have a pending request. The event counts also when one core is halted and the other is not halted.<br><br>The event can count occurrences for this core or both cores.   |
| 3AH       | 00H   | EIST_TRANS  | Number of Enhanced Intel SpeedStep Technology (EIST) transitions. | This event counts the number of transitions that include a frequency change, either with or without voltage change. This includes Enhanced Intel SpeedStep Technology (EIST) and TM2 transitions.<br><br>The event is incremented only while the counting core is in C0 state. Since transitions to higher-numbered CxE states and TM2 transitions include a frequency change or voltage transition, the event is incremented accordingly.  |



**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value     | Event Name                        | Definition  | Description and Comment   |
|-----------|-----------------|-----------------------------------|---|---|
| 3BH       | COH             | THERMAL_TRIP                      | Number of thermal trips.                                | This event counts the number of thermal trips. A thermal trip occurs whenever the processor temperature exceeds the thermal trip threshold temperature.<br><br>Following a thermal trip, the processor automatically reduces frequency and voltage. The processor checks the temperature every millisecond and returns to normal when the temperature falls below the thermal trip threshold temperature.   |
| 3CH       | 00H             | CPU_CLK_UNHALTED.CORE_P           | Core cycles when core is not halted.                    | This event counts the number of core cycles while the core is not in a halt state. The core enters the halt state when it is running the HLT instruction. This event is a component in many key event ratios.<br><br>The core frequency may change due to transitions associated with Enhanced Intel SpeedStep Technology or TM2. For this reason, this event may have a changing ratio in regard to time.<br><br>When the core frequency is constant, this event can give approximate elapsed time while the core not in halt state. This is an architectural performance event. |
| 3CH       | 01H             | CPU_CLK_UNHALTED.BUS              | Bus cycles when core is not halted.                     | This event counts the number of bus cycles while the core is not in the halt state. This event can give a measurement of the elapsed time while the core was not in the halt state. The core enters the halt state when it is running the HLT instruction.<br><br>The event also has a constant ratio with CPU_CLK_UNHALTED.REF event, which is the maximum bus to processor frequency ratio.<br><br>Non-halted bus cycles are a component in many key event ratios.  |
| 3CH       | 02H             | CPU_CLK_UNHALTED.NO_OTHER         | Bus cycles when core is active and the other is halted. | This event counts the number of bus cycles during which the core remains non-halted and the other core on the processor is halted.<br><br>This event can be used to determine the amount of parallelism exploited by an application or a system. Divide this event count by the bus frequency to determine the amount of time that only one core was in use.  |
| 40H       | See Table 18-64 | L1D_CACHE_LD.(Cache Line State)   | L1 cacheable data reads.                                | This event counts the number of data reads from cacheable memory. Locked reads are not counted.   |
| 41H       | See Table 18-64 | L1D_CACHE_ST.(Cache Line State)   | L1 cacheable data writes.                               | This event counts the number of data writes to cacheable memory. Locked writes are not counted.   |
| 42H       | See Table 18-64 | L1D_CACHE_LOCK.(Cache Line State) | L1 data cacheable locked reads.                         | This event counts the number of locked data reads from cacheable memory.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name              | Definition  | Description and Comment   |
|-----------|-------------|-------------------------|---|---|
| 42H       | 10H         | L1D_CACHE_LOCK_DURATION | Duration of L1 data cacheable locked operation.                                     | This event counts the number of cycles during which any cache line is locked by any locking instruction. Locking happens at retirement and therefore the event does not occur for instructions that are speculatively executed. Locking duration is shorter than locked instruction execution duration.   |
| 43H       | 01H         | L1D_ALL_REF             | All references to the L1 data cache.  | This event counts all references to the L1 data cache, including all loads and stores with any memory types. The event counts memory accesses only when they are actually performed. For example, a load blocked by unknown store address and later performed is only counted once. The event includes non-cacheable accesses, such as I/O accesses.  |
| 43H       | 02H         | L1D_ALL_CACHE_REF       | L1 Data cacheable reads and writes.   | This event counts the number of data reads and writes from cacheable memory, including locked operations. This event is a sum of: <ul style="list-style-type: none"> <li>▪ L1D_CACHE_LD.MESI</li> <li>▪ L1D_CACHE_ST.MESI</li> <li>▪ L1D_CACHE_LOCK.MESI</li> </ul>   |
| 45H       | 0FH         | L1D_REPL                | Cache lines allocated in the L1 data cache.   | This event counts the number of lines brought into the L1 data cache.   |
| 46H       | 00H         | L1D_M_REPL              | Modified cache lines allocated in the L1 data cache.                                | This event counts the number of modified lines brought into the L1 data cache.  |
| 47H       | 00H         | L1D_M_EVICT             | Modified cache lines evicted from the L1 data cache.                                | This event counts the number of modified lines evicted from the L1 data cache, whether due to replacement or by snoop HITM intervention.  |
| 48H       | 00H         | L1D_PEND_MISS           | Total number of outstanding L1 data cache misses at any cycle.                      | This event counts the number of outstanding L1 data cache misses at any cycle. An L1 data cache miss is outstanding from the cycle on which the miss is determined until the first chunk of data is available. This event counts: <ul style="list-style-type: none"> <li>▪ All cacheable demand requests.</li> <li>▪ L1 data cache hardware prefetch requests.</li> <li>▪ Requests to write through memory.</li> <li>▪ Requests to write combine memory.</li> </ul> Uncacheable requests are not counted. The count of this event divided by the number of L1 data cache misses, L1D_REPL, is the average duration in core cycles of an L1 data cache miss. |
| 49H       | 01H         | L1D_SPLIT.LOADS         | Cache line split loads from the L1 data cache.                                      | This event counts the number of load operations that span two cache lines. Such load operations are also called split loads. Split load operations are executed at retirement.  |
| 49H       | 02H         | L1D_SPLIT.STORES        | Cache line split stores to the L1 data cache.                                       | This event counts the number of store operations that span two cache lines.   |
| 4BH       | 00H         | SSE_PRE_MISS.NTA        | Streaming SIMD Extensions (SSE) Prefetch NTA instructions missing all cache levels. | This event counts the number of times the SSE instructions prefetchNTA were executed and missed all cache levels. Due to speculation an executed instruction might not retire. This instruction prefetches the data to the L1 data cache.   |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value                      | Event Name                                     | Definition   | Description and Comment   |
|-----------|----------------------------------|--|--|---|
| 4BH       | 01H                              | SSE_PRE_MISS.L1                                | Streaming SIMD Extensions (SSE) PrefetchT0 instructions missing all cache levels.                | This event counts the number of times the SSE instructions prefetchT0 were executed and missed all cache levels. Due to speculation executed instruction might not retire. The prefetchT0 instruction prefetches data to the L2 cache and L1 data cache.  |
| 4BH       | 02H                              | SSE_PRE_MISS.L2                                | Streaming SIMD Extensions (SSE) PrefetchT1 and PrefetchT2 instructions missing all cache levels. | This event counts the number of times the SSE instructions prefetchT1 and prefetchT2 were executed and missed all cache levels. Due to speculation, an executed instruction might not retire. The prefetchT1 and PrefetchNT2 instructions prefetch data to the L2 cache.  |
| 4CH       | 00H                              | LOAD_HIT_PRE                                   | Load operations conflicting with a software prefetch to the same address.                        | This event counts load operations sent to the L1 data cache while a previous Streaming SIMD Extensions (SSE) prefetch instruction to the same cache line has started prefetching but has not yet finished.  |
| 4EH       | 10H                              | L1D_PREFETCH.REQUESTS                          | L1 data cache prefetch requests.   | This event counts the number of times the L1 data cache requested to prefetch a data cache line. Requests can be rejected when the L2 cache is busy and resubmitted later or lost. All requests are counted, including those that are rejected.   |
| 60H       | See Table 18-61 and Table 18-62. | BUS_REQUEST_OUTSTANDING. (Core and Bus Agents) | Outstanding cacheable data read bus requests duration.   | This event counts the number of pending full cache line read transactions on the bus occurring in each cycle. A read transaction is pending from the cycle it is sent on the bus until the full cache line is received by the processor. The event counts only full-line cacheable read requests from either the L1 data cache or the L2 prefetchers. It does not count Read for Ownership transactions, instruction byte fetch transactions, or any other bus transaction.   |
| 61H       | See Table 18-62.                 | BUS_BNR_DRV. (Bus Agents)                      | Number of Bus Not Ready signals asserted.  | This event counts the number of Bus Not Ready (BNR) signals that the processor asserts on the bus to suspend additional bus requests by other bus agents. A bus agent asserts the BNR signal when the number of data and snoop transactions is close to the maximum that the bus can handle. To obtain the number of bus cycles during which the BNR signal is asserted, multiply the event count by two. While this signal is asserted, new transactions cannot be submitted on the bus. As a result, transaction latency may have higher impact on program performance.   |
| 62H       | See Table 18-62.                 | BUS_DRDY_CLOCKS.(Bus Agents)                   | Bus cycles when data is sent on the bus.   | This event counts the number of bus cycles during which the DRDY (Data Ready) signal is asserted on the bus. The DRDY signal is asserted when data is sent on the bus. With the 'THIS_AGENT' mask this event counts the number of bus cycles during which this agent (the processor) writes data on the bus back to memory or to other bus agents. This includes all explicit and implicit data writebacks, as well as partial writes. With the 'ALL_AGENTS' mask, this event counts the number of bus cycles during which any bus agent sends data on the bus. This includes all data reads and writes on the bus. |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value                      | Event Name                              | Definition                                | Description and Comment   |
|-----------|----------------------------------|---|---|---|
| 63H       | See Table 18-61 and Table 18-62. | BUS_LOCK_CLOCKS.(Core and Bus Agents)   | Bus cycles when a LOCK signal asserted.   | This event counts the number of bus cycles, during which the LOCK signal is asserted on the bus. A LOCK signal is asserted when there is a locked memory access, due to: <ul style="list-style-type: none"> <li>Uncacheable memory.</li> <li>Locked operation that spans two cache lines.</li> <li>Page-walk from an uncacheable page table.</li> </ul> Bus locks have a very high performance penalty and it is highly recommended to avoid such accesses. |
| 64H       | See Table 18-61.                 | BUS_DATA_RCV.(Core)                     | Bus cycles while processor receives data. | This event counts the number of bus cycles during which the processor is busy receiving data.   |
| 65H       | See Table 18-61 and Table 18-62. | BUS_TRANS_BRD.(Core and Bus Agents)     | Burst read bus transactions.              | This event counts the number of burst read transactions including: <ul style="list-style-type: none"> <li>L1 data cache read misses (and L1 data cache hardware prefetches).</li> <li>L2 hardware prefetches by the DPL and L2 streamer.</li> <li>IFU read misses of cacheable lines.</li> </ul> It does not include RFO transactions.  |
| 66H       | See Table 18-61 and Table 18-62. | BUS_TRANS_RFO.(Core and Bus Agents)     | RFO bus transactions.                     | This event counts the number of Read For Ownership (RFO) bus transactions, due to store operations that miss the L1 data cache and the L2 cache. It also counts RFO bus transactions due to locked operations.  |
| 67H       | See Table 18-61 and Table 18-62. | BUS_TRANS_WB.(Core and Bus Agents)      | Explicit writeback bus transactions.      | This event counts all explicit writeback bus transactions due to dirty line evictions. It does not count implicit writebacks due to invalidation by a snoop request.  |
| 68H       | See Table 18-61 and Table 18-62. | BUS_TRANS_IFETCH.(Core and Bus Agents)  | Instruction-fetch bus transactions.       | This event counts all instruction fetch full cache line bus transactions.   |
| 69H       | See Table 18-61 and Table 18-62. | BUS_TRANS_INVALID.(Core and Bus Agents) | Invalidate bus transactions.              | This event counts all invalidate transactions. Invalidate transactions are generated when: <ul style="list-style-type: none"> <li>A store operation hits a shared line in the L2 cache.</li> <li>A full cache line write misses the L2 cache or hits a shared line in the L2 cache.</li> </ul>  |
| 6AH       | See Table 18-61 and Table 18-62. | BUS_TRANS_PWR.(Core and Bus Agents)     | Partial write bus transaction.            | This event counts partial write bus transactions.   |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value                      | Event Name                             | Definition                                | Description and Comment   |
|-----------|----------------------------------|--|---|---|
| 6BH       | See Table 18-61 and Table 18-62. | BUS_TRANS_P.(Core and Bus Agents)      | Partial bus transactions.                 | This event counts all (read and write) partial bus transactions.  |
| 6CH       | See Table 18-61 and Table 18-62. | BUS_TRANS_IO.(Core and Bus Agents)     | IO bus transactions.                      | This event counts the number of completed I/O bus transactions as a result of IN and OUT instructions. The count does not include memory mapped IO.   |
| 6DH       | See Table 18-61 and Table 18-62. | BUS_TRANS_DEF.(Core and Bus Agents)    | Deferred bus transactions.                | This event counts the number of deferred transactions.  |
| 6EH       | See Table 18-61 and Table 18-62. | BUS_TRANS_BURST.(Core and Bus Agents)  | Burst (full cache-line) bus transactions. | This event counts burst (full cache line) transactions including: <ul style="list-style-type: none"> <li>▪ Burst reads.</li> <li>▪ RFOs.</li> <li>▪ Explicit writebacks.</li> <li>▪ Write combine lines.</li> </ul>   |
| 6FH       | See Table 18-61 and Table 18-62. | BUS_TRANS_MEM.(Core and Bus Agents)    | Memory bus transactions.                  | This event counts all memory bus transactions including: <ul style="list-style-type: none"> <li>▪ Burst transactions.</li> <li>▪ Partial reads and writes - invalidate transactions.</li> </ul> The BUS_TRANS_MEM count is the sum of BUS_TRANS_BURST, BUS_TRANS_P and BUS_TRANS_IVAL.  |
| 70H       | See Table 18-61 and Table 18-62. | BUS_TRANS_ANY.(Core and Bus Agents)    | All bus transactions.                     | This event counts all bus transactions. This includes: <ul style="list-style-type: none"> <li>▪ Memory transactions.</li> <li>▪ IO transactions (non memory-mapped).</li> <li>▪ Deferred transaction completion.</li> <li>▪ Other less frequent transactions, such as interrupts.</li> </ul>  |
| 77H       | See Table 18-61 and Table 18-65. | EXT_SNOOP.(Bus Agents, Snoop Response) | External snoops.                          | This event counts the snoop responses to bus transactions. Responses can be counted separately by type and by bus agent.<br><br>With the 'THIS_AGENT' mask, the event counts snoop responses from this processor to bus transactions sent by this processor. With the 'ALL_AGENTS' mask the event counts all snoop responses seen on the bus. |
| 78H       | See Table 18-61 and Table 18-66. | CMP_SNOOP.(Core, Snoop Type)           | L1 data cache snooped by other core.      | This event counts the number of times the L1 data cache is snooped for a cache line that is needed by the other core in the same processor. The cache line is either missing in the L1 instruction or data caches of the other core, or is available for reading only and the other core wishes to write the cache line.                      |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value                      | Event Name                                | Definition                            | Description and Comment  |
|-----------|----------------------------------|---|---------------------------------------|--|
|           |                                  |   |                                       | <p>The snoop operation may change the cache line state. If the other core issued a read request that hit this core in E state, typically the state changes to S state in this core. If the other core issued a read for ownership request (due a write miss or hit to S state) that hits this core's cache line in E or S state, this typically results in invalidation of the cache line in this core. If the snoop hits a line in M state, the state is changed at a later opportunity.</p> <p>These snoops are performed through the L1 data cache store port. Therefore, frequent snoops may conflict with extensive stores to the L1 data cache, which may increase store latency and impact performance.</p> |
| 7AH       | See Table 18-62.                 | BUS_HIT_DRV.<br>(Bus Agents)              | HIT signal asserted.                  | This event counts the number of bus cycles during which the processor drives the HIT# pin to signal HIT snoop response.  |
| 7BH       | See Table 18-62.                 | BUS_HITM_DRV.<br>(Bus Agents)             | HITM signal asserted.                 | This event counts the number of bus cycles during which the processor drives the HITM# pin to signal HITM snoop response.  |
| 7DH       | See Table 18-61.                 | BUSQ_EMPTY.<br>(Core)                     | Bus queue empty.                      | <p>This event counts the number of cycles during which the core did not have any pending transactions in the bus queue. It also counts when the core is halted and the other core is not halted.</p> <p>This event can count occurrences for this core or both cores.</p>  |
| 7EH       | See Table 18-61 and Table 18-62. | SNOOP_STALL_DRV.<br>(Core and Bus Agents) | Bus stalled for snoops.               | <p>This event counts the number of times that the bus snoop stall signal is asserted. To obtain the number of bus cycles during which snoops on the bus are prohibited, multiply the event count by two.</p> <p>During the snoop stall cycles, no new bus transactions requiring a snoop response can be initiated on the bus. A bus agent asserts a snoop stall signal if it cannot response to a snoop request within three bus cycles.</p>  |
| 7FH       | See Table 18-61.                 | BUS_IO_WAIT.<br>(Core)                    | IO requests waiting in the bus queue. | <p>This event counts the number of core cycles during which IO requests wait in the bus queue. With the SELF modifier this event counts IO requests per core.</p> <p>With the BOTH_CORE modifier, this event increments by one for any cycle for which there is a request from either core.</p>  |
| 80H       | 00H                              | L1I_READS                                 | Instruction fetches.                  | This event counts all instruction fetches, including uncacheable fetches that bypass the Instruction Fetch Unit (IFU).   |
| 81H       | 00H                              | L1I_MISSES                                | Instruction Fetch Unit misses.        | <p>This event counts all instruction fetches that miss the Instruction Fetch Unit (IFU) or produce memory requests. This includes uncacheable fetches.</p> <p>An instruction fetch miss is counted only once and not once for every cycle it is outstanding.</p>   |
| 82H       | 02H                              | ITLB.SMALL_MISS                           | ITLB small page misses.               | This event counts the number of instruction fetches from small pages that miss the ITLB.   |
| 82H       | 10H                              | ITLB.LARGE_MISS                           | ITLB large page misses.               | This event counts the number of instruction fetches from large pages that miss the ITLB.   |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name            | Definition   | Description and Comment   |
|-----------|-------------|-----------------------|--|---|
| 82H       | 40H         | ITLB.FLUSH            | ITLB flushes.  | This event counts the number of ITLB flushes. This usually happens upon CR3 or CR0 writes, which are executed by the operating system during process switches.  |
| 82H       | 12H         | ITLB.MISSES           | ITLB misses.   | This event counts the number of instruction fetches from either small or large pages that miss the ITLB.  |
| 83H       | 02H         | INST_QUEUE.FULL       | Cycles during which the instruction queue is full.                       | This event counts the number of cycles during which the instruction queue is full. In this situation, the core front end stops fetching more instructions. This is an indication of very long stalls in the back-end pipeline stages.   |
| 86H       | 00H         | CYCLES_L1_MEM_STALLED | Cycles during which instruction fetches stalled.                         | This event counts the number of cycles for which an instruction fetch stalls, including stalls due to any of the following reasons: <ul style="list-style-type: none"> <li>▪ Instruction Fetch Unit cache misses.</li> <li>▪ Instruction TLB misses.</li> <li>▪ Instruction TLB faults.</li> </ul>  |
| 87H       | 00H         | ILD_STALL             | Instruction Length Decoder stall cycles due to a length changing prefix. | This event counts the number of cycles during which the instruction length decoder uses the slow length decoder. Usually, instruction length decoding is done in one cycle. When the slow decoder is used, instruction decoding requires 6 cycles.<br><br>The slow decoder is used in the following cases: <ul style="list-style-type: none"> <li>▪ Operand override prefix (66H) preceding an instruction with immediate data.</li> <li>▪ Address override prefix (67H) preceding an instruction with a modr/m in real, big real, 16-bit protected or 32-bit protected modes.</li> </ul> To avoid instruction length decoding stalls, generate code using imm8 or imm32 values instead of imm16 values. If you must use an imm16 value, store the value in a register using "mov reg, imm32" and use the register format of the instruction. |
| 88H       | 00H         | BR_INST_EXEC          | Branch instructions executed.  | This event counts all executed branches (not necessarily retired). This includes only instructions and not micro-op branches.<br><br>Frequent branching is not necessarily a major performance issue. However frequent branch mispredictions may be a problem.  |
| 89H       | 00H         | BR_MISSP_EXEC         | Mispredicted branch instructions executed.                               | This event counts the number of mispredicted branch instructions that were executed.  |
| 8AH       | 00H         | BR_BAC_MISSP_EXEC     | Branch instructions mispredicted at decoding.                            | This event counts the number of branch instructions that were mispredicted at decoding.   |
| 8BH       | 00H         | BR_CND_EXEC           | Conditional branch instructions executed.                                | This event counts the number of conditional branch instructions executed, but not necessarily retired.  |
| 8CH       | 00H         | BR_CND_MISSP_EXEC     | Mispredicted conditional branch instructions executed.                   | This event counts the number of mispredicted conditional branch instructions that were executed.  |
| 8DH       | 00H         | BR_IND_EXEC           | Indirect branch instructions executed.                                   | This event counts the number of indirect branch instructions that were executed.  |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name               | Definition   | Description and Comment  |
|-----------|-------------|--------------------------|--|--|
| 8EH       | 00H         | BR_IND_MISSP_EXEC        | Mispredicted indirect branch instructions executed.  | This event counts the number of mispredicted indirect branch instructions that were executed.  |
| 8FH       | 00H         | BR_RET_EXEC              | RET instructions executed.                           | This event counts the number of RET instructions that were executed.   |
| 90H       | 00H         | BR_RET_MISSP_EXEC        | Mispredicted RET instructions executed.              | This event counts the number of mispredicted RET instructions that were executed.  |
| 91H       | 00H         | BR_RET_BAC_MISSP_EXEC    | RET instructions executed mispredicted at decoding.  | This event counts the number of RET instructions that were executed and were mispredicted at decoding.   |
| 92H       | 00H         | BR_CALL_EXEC             | CALL instructions executed.                          | This event counts the number of CALL instructions executed.  |
| 93H       | 00H         | BR_CALL_MISSP_EXEC       | Mispredicted CALL instructions executed.             | This event counts the number of mispredicted CALL instructions that were executed.   |
| 94H       | 00H         | BR_IND_CALL_EXEC         | Indirect CALL instructions executed.                 | This event counts the number of indirect CALL instructions that were executed.   |
| 97H       | 00H         | BR_TKN_BUBBLE_1          | Branch predicted taken with bubble 1.                | The events BR_TKN_BUBBLE_1 and BR_TKN_BUBBLE_2 together count the number of times a taken branch prediction incurred a one-cycle penalty. The penalty incurs when: <ul style="list-style-type: none"> <li>Too many taken branches are placed together. To avoid this, unroll loops and add a non-taken branch in the middle of the taken sequence.</li> <li>The branch target is unaligned. To avoid this, align the branch target.</li> </ul> |
| 98H       | 00H         | BR_TKN_BUBBLE_2          | Branch predicted taken with bubble 2.                | The events BR_TKN_BUBBLE_1 and BR_TKN_BUBBLE_2 together count the number of times a taken branch prediction incurred a one-cycle penalty. The penalty incurs when: <ul style="list-style-type: none"> <li>Too many taken branches are placed together. To avoid this, unroll loops and add a non-taken branch in the middle of the taken sequence.</li> <li>The branch target is unaligned. To avoid this, align the branch target.</li> </ul> |
| A0H       | 00H         | RS_UOPS_DISPATCHED       | Micro-ops dispatched for execution.                  | This event counts the number of micro-ops dispatched for execution. Up to six micro-ops can be dispatched in each cycle.   |
| A1H       | 01H         | RS_UOPS_DISPATCHED.PORT0 | Cycles micro-ops dispatched for execution on port 0. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Issue Ports are described in <i>Intel® 64 and IA-32 Architectures Optimization Reference Manual</i> . Use IA32_PMC0 only.   |
| A1H       | 02H         | RS_UOPS_DISPATCHED.PORT1 | Cycles micro-ops dispatched for execution on port 1. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Use IA32_PMC0 only.   |
| A1H       | 04H         | RS_UOPS_DISPATCHED.PORT2 | Cycles micro-ops dispatched for execution on port 2. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Use IA32_PMC0 only.   |



**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                 | Definition   | Description and Comment  |
|-----------|-------------|----------------------------|--|--|
| A1H       | 08H         | RS_UOPS_DISPATCHED.PORT3   | Cycles micro-ops dispatched for execution on port 3. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Use IA32_PMC0 only.   |
| A1H       | 10H         | RS_UOPS_DISPATCHED.PORT4   | Cycles micro-ops dispatched for execution on port 4. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Use IA32_PMC0 only.   |
| A1H       | 20H         | RS_UOPS_DISPATCHED.PORT5   | Cycles micro-ops dispatched for execution on port 5. | This event counts the number of cycles for which micro-ops dispatched for execution. Each cycle, at most one micro-op can be dispatched on the port. Use IA32_PMC0 only.   |
| AAH       | 01H         | MACRO_INSTS_DECODED        | Instructions decoded.                                | This event counts the number of instructions decoded (but not necessarily executed or retired).  |
| AAH       | 08H         | MACRO_INSTS_CISC_DECODED   | CISC Instructions decoded.                           | This event counts the number of complex instructions decoded. Complex instructions usually have more than four micro-ops. Only one complex instruction can be decoded at a time.   |
| ABH       | 01H         | ESP.SYNCH                  | ESP register content synchronization.                | This event counts the number of times that the ESP register is explicitly used in the address expression of a load or store operation, after it is implicitly used, for example by a push or a pop instruction.<br><br>ESP synch micro-op uses resources from the rename pipeline and up to retirement. The expected ratio of this event divided by the number of ESP implicit changes is 0.2. If the ratio is higher, consider rearranging your code to avoid ESP synchronization events. |
| ABH       | 02H         | ESP.ADDITIONS              | ESP register automatic additions.                    | This event counts the number of ESP additions performed automatically by the decoder. A high count of this event is good, since each automatic addition performed by the decoder saves a micro-op from the execution units.<br><br>To maximize the number of ESP additions performed automatically by the decoder, choose instructions that implicitly use the ESP, such as PUSH, POP, CALL, and RET instructions whenever possible.   |
| B0H       | 00H         | SIMD_UOPS_EXEC             | SIMD micro-ops executed (excluding stores).          | This event counts all the SIMD micro-ops executed. It does not count MOVQ and MOVD stores from register to memory.   |
| B1H       | 00H         | SIMD_SAT_UOP_EXEC          | SIMD saturated arithmetic micro-ops executed.        | This event counts the number of SIMD saturated arithmetic micro-ops executed.  |
| B3H       | 01H         | SIMD_UOP_TYPE_EXEC.MUL     | SIMD packed multiply micro-ops executed.             | This event counts the number of SIMD packed multiply micro-ops executed.   |
| B3H       | 02H         | SIMD_UOP_TYPE_EXEC.SHIFT   | SIMD packed shift micro-ops executed.                | This event counts the number of SIMD packed shift micro-ops executed.  |
| B3H       | 04H         | SIMD_UOP_TYPE_EXEC.PACK    | SIMD pack micro-ops executed.                        | This event counts the number of SIMD pack micro-ops executed.  |
| B3H       | 08H         | SIMD_UOP_TYPE_EXEC.UNPACK  | SIMD unpack micro-ops executed.                      | This event counts the number of SIMD unpack micro-ops executed.  |
| B3H       | 10H         | SIMD_UOP_TYPE_EXEC.LOGICAL | SIMD packed logical micro-ops executed.              | This event counts the number of SIMD packed logical micro-ops executed.  |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name                    | Definition   | Description and Comment  |
|-----------|-------------|-------------------------------|--|--|
| B3H       | 20H         | SIMD_UOP_TYPE_EXEC.ARITHMETIC | SIMD packed arithmetic micro-ops executed.                       | This event counts the number of SIMD packed arithmetic micro-ops executed.   |
| COH       | 00H         | INST_RETIRED.ANY_P            | Instructions retired.  | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers. INST_RETIRED.ANY_P is an architectural performance event.  |
| COH       | 01H         | INST_RETIRED.LOADS            | Instructions retired, which contain a load.                      | This event counts the number of instructions retired that contain a load operation.  |
| COH       | 02H         | INST_RETIRED.STORES           | Instructions retired, which contain a store.                     | This event counts the number of instructions retired that contain a store operation.   |
| COH       | 04H         | INST_RETIRED.OTHER            | Instructions retired, with no load or store operation.           | This event counts the number of instructions retired that do not contain a load or a store operation.  |
| C1H       | 01H         | X87_OPS_RETIRED.FXCH          | FXCH instructions retired.                                       | This event counts the number of FXCH instructions retired. Modern compilers generate more efficient code and are less likely to use this instruction. If you obtain a high count for this event consider recompiling the code.   |
| C1H       | FEH         | X87_OPS_RETIRED.ANY           | Retired floating-point computational operations (precise event). | <p>This event counts the number of floating-point computational operations retired. It counts:</p> <ul style="list-style-type: none"> <li>▪ Floating point computational operations executed by the assist handler.</li> <li>▪ Sub-operations of complex floating-point instructions like transcendental instructions.</li> </ul> <p>This event does not count:</p> <ul style="list-style-type: none"> <li>▪ Floating-point computational operations that cause traps or assists.</li> <li>▪ Floating-point loads and stores.</li> </ul> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p>  |
| C2H       | 01H         | UOPS_RETIRED.LD_IND_BR        | Fused load+op or load+indirect branch retired.                   | <p>This event counts the number of retired micro-ops that fused a load with another operation. This includes:</p> <ul style="list-style-type: none"> <li>▪ Fusion of a load and an arithmetic operation, such as with the following instruction: ADD EAX, [EBX] where the content of the memory location specified by EBX register is loaded, added to EXA register, and the result is stored in EAX.</li> <li>▪ Fusion of a load and a branch in an indirect branch operation, such as with the following instructions: <ul style="list-style-type: none"> <li>▪ JMP [RDI+200]</li> <li>▪ RET</li> </ul> </li> <li>▪ Fusion decreases the number of micro-ops in the processor pipeline. A high value for this event count indicates that the code is using the processor resources effectively.</li> </ul> |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                     | Definition   | Description and Comment   |
|-----------|-------------|--------------------------------|--|---|
| C2H       | 02H         | UOPS_RETIREDD.<br>STD_STA      | Fused store address + data retired.                | This event counts the number of store address calculations that are fused with store data emission into one micro-op. Traditionally, each store operation required two micro-ops. This event counts fusion of retired micro-ops only. Fusion decreases the number of micro-ops in the processor pipeline. A high value for this event count indicates that the code is using the processor resources effectively.   |
| C2H       | 04H         | UOPS_RETIREDD.<br>MACRO_FUSION | Retired instruction pairs fused into one micro-op. | This event counts the number of times CMP or TEST instructions were fused with a conditional branch instruction into one micro-op. It counts fusion by retired micro-ops only. Fusion decreases the number of micro-ops in the processor pipeline. A high value for this event count indicates that the code uses the processor resources more effectively.   |
| C2H       | 07H         | UOPS_RETIREDD.<br>FUSED        | Fused micro-ops retired.                           | This event counts the total number of retired fused micro-ops. The counts include the following fusion types: <ul style="list-style-type: none"> <li>▪ Fusion of load operation with an arithmetic operation or with an indirect branch (counted by event UOPS_RETIREDD.LD_IND_BR)</li> <li>▪ Fusion of store address and data (counted by event UOPS_RETIREDD.STD_STA)</li> <li>▪ Fusion of CMP or TEST instruction with a conditional branch instruction (counted by event UOPS_RETIREDD.MACRO_FUSION)</li> </ul> Fusion decreases the number of micro-ops in the processor pipeline. A high value for this event count indicates that the code is using the processor resources effectively. |
| C2H       | 08H         | UOPS_RETIREDD.<br>NON_FUSED    | Non-fused micro-ops retired.                       | This event counts the number of micro-ops retired that were not fused.  |
| C2H       | 0FH         | UOPS_RETIREDD.<br>ANY          | Micro-ops retired.                                 | This event counts the number of micro-ops retired. The processor decodes complex macro instructions into a sequence of simpler micro-ops. Most instructions are composed of one or two micro-ops. Some instructions are decoded into longer sequences such as repeat instructions, floating point transcendental instructions, and assists. In some cases micro-op sequences are fused or whole instructions are fused into one micro-op. See other UOPS_RETIREDD events for differentiating retired fused and non-fused micro-ops.   |
| C3H       | 01H         | MACHINE_NUKES.SMC              | Self-Modifying Code detected.                      | This event counts the number of times that a program writes to a code section. Self-modifying code causes a severe penalty in all Intel 64 and IA-32 processors.  |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name                        | Definition   | Description and Comment   |
|-----------|-------------|-----------------------------------|--|---|
| C3H       | 04H         | MACHINE_NUKES.MEM_ORDER           | Execution pipeline restart due to memory ordering conflict or memory disambiguation misprediction. | This event counts the number of times the pipeline is restarted due to either multi-threaded memory ordering conflicts or memory disambiguation misprediction.<br><br>A multi-threaded memory ordering conflict occurs when a store, which is executed in another core, hits a load that is executed out of order in this core but not yet retired. As a result, the load needs to be restarted to satisfy the memory ordering model.<br><br>See Chapter 8, "Multiple-Processor Management" in the <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A</i> .<br><br>To count memory disambiguation mispredictions, use the event MEMORY_DISAMBIGUATION.RESET. |
| C4H       | 00H         | BR_INST_RETIRED.ANY               | Retired branch instructions.   | This event counts the number of branch instructions retired. This is an architectural performance event.  |
| C4H       | 01H         | BR_INST_RETIRED.PRED_NOT_TAKEN    | Retired branch instructions that were predicted not-taken.   | This event counts the number of branch instructions retired that were correctly predicted to be not-taken.  |
| C4H       | 02H         | BR_INST_RETIRED.MISPRED_NOT_TAKEN | Retired branch instructions that were mispredicted not-taken.                                      | This event counts the number of branch instructions retired that were mispredicted and not-taken.   |
| C4H       | 04H         | BR_INST_RETIRED.PRED_TAKEN        | Retired branch instructions that were predicted taken.   | This event counts the number of branch instructions retired that were correctly predicted to be taken.  |
| C4H       | 08H         | BR_INST_RETIRED.MISPRED_TAKEN     | Retired branch instructions that were mispredicted taken.  | This event counts the number of branch instructions retired that were mispredicted and taken.   |
| C4H       | 0CH         | BR_INST_RETIRED.TAKEN             | Retired taken branch instructions.   | This event counts the number of branches retired that were taken.   |
| C5H       | 00H         | BR_INST_RETIRED.MISPRED           | Retired mispredicted branch instructions. (precise event)  | This event counts the number of retired branch instructions that were mispredicted by the processor. A branch misprediction occurs when the processor predicts that the branch would be taken, but it is not, or vice-versa.<br><br>This is an architectural performance event.   |
| C6H       | 01H         | CYCLES_INT_MASKED                 | Cycles during which interrupts are disabled.   | This event counts the number of cycles during which interrupts are disabled.  |
| C6H       | 02H         | CYCLES_INT_PENDING_AND_MASKED     | Cycles during which interrupts are pending and disabled.   | This event counts the number of cycles during which there are pending interrupts but interrupts are disabled.   |
| C7H       | 01H         | SIMD_INST_RETIRED.PACKED_SINGLE   | Retired SSE packed-single instructions.  | This event counts the number of SSE packed-single instructions retired.   |
| C7H       | 02H         | SIMD_INST_RETIRED.SCALAR_SINGLE   | Retired SSE scalar-single instructions.  | This event counts the number of SSE scalar-single instructions retired.   |
| C7H       | 04H         | SIMD_INST_RETIRED.PACKED_DOUBLE   | Retired SSE2 packed-double instructions.   | This event counts the number of SSE2 packed-double instructions retired.  |
| C7H       | 08H         | SIMD_INST_RETIRED.SCALAR_DOUBLE   | Retired SSE2 scalar-double instructions.   | This event counts the number of SSE2 scalar-double instructions retired.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                            | Definition   | Description and Comment  |
|-----------|-------------|---------------------------------------|--|--|
| C7H       | 10H         | SIMD_INST_RETIRE.D.VECTOR             | Retired SSE2 vector integer instructions.              | This event counts the number of SSE2 vector integer instructions retired.  |
| C7H       | 1FH         | SIMD_INST_RETIRE.ANY                  | Retired Streaming SIMD instructions (precise event).   | This event counts the overall number of retired SIMD instructions that use XMM registers. To count each type of SIMD instruction separately, use the following events: <ul style="list-style-type: none"> <li>▪ SIMD_INST_RETIRE.PACKED_SINGLE</li> <li>▪ SIMD_INST_RETIRE.SCALAR_SINGLE</li> <li>▪ SIMD_INST_RETIRE.PACKED_DOUBLE</li> <li>▪ SIMD_INST_RETIRE.SCALAR_DOUBLE</li> <li>▪ and SIMD_INST_RETIRE.VECTOR</li> </ul> When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event. |
| C8H       | 00H         | HW_INT_RCV                            | Hardware interrupts received.                          | This event counts the number of hardware interrupts received by the processor.   |
| C9H       | 00H         | ITLB_MISS_RETIRE                      | Retired instructions that missed the ITLB.             | This event counts the number of retired instructions that missed the ITLB when they were fetched.  |
| CAH       | 01H         | SIMD_COMP_INST_RETIRE.PACKED_SINGLE   | Retired computational SSE packed-single instructions.  | This event counts the number of computational SSE packed-single instructions retired. Computational instructions perform arithmetic computations (for example: add, multiply and divide).<br>Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.   |
| CAH       | 02H         | SIMD_COMP_INST_RETIRE.SCALAR_SINGLE   | Retired computational SSE scalar-single instructions.  | This event counts the number of computational SSE scalar-single instructions retired. Computational instructions perform arithmetic computations (for example: add, multiply and divide).<br>Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.   |
| CAH       | 04H         | SIMD_COMP_INST_RETIRE.PACKED_DOUBLE   | Retired computational SSE2 packed-double instructions. | This event counts the number of computational SSE2 packed-double instructions retired. Computational instructions perform arithmetic computations (for example: add, multiply and divide).<br>Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.  |
| CAH       | 08H         | SIMD_COMP_INST_RETIRE.D.SCALAR_DOUBLE | Retired computational SSE2 scalar-double instructions. | This event counts the number of computational SSE2 scalar-double instructions retired. Computational instructions perform arithmetic computations (for example: add, multiply and divide).<br>Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                      | Definition  | Description and Comment  |
|-----------|-------------|---------------------------------|---|--|
| CBH       | 01H         | MEM_LOAD_RETIREDD.L1D_MISS      | Retired loads that miss the L1 data cache (precise event).  | <p>This event counts the number of retired load operations that missed the L1 data cache. This includes loads from cache lines that are currently being fetched, due to a previous L1 data cache miss to the same cache line.</p> <p>This event counts loads from cacheable memory only. The event does not count loads by software prefetches.</p> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p> <p>Use IA32_PMC0 only.</p>  |
| CBH       | 02H         | MEM_LOAD_RETIREDD.L1D_LINE_MISS | L1 data cache line missed by retired loads (precise event). | <p>This event counts the number of load operations that miss the L1 data cache and send a request to the L2 cache to fetch the missing cache line. That is the missing cache line fetching has not yet started.</p> <p>The event count is equal to the number of cache lines fetched from the L2 cache by retired loads.</p> <p>This event counts loads from cacheable memory only. The event does not count loads by software prefetches.</p> <p>The event might not be counted if the load is blocked (see LOAD_BLOCK events).</p> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p> <p>Use IA32_PMC0 only.</p> |
| CBH       | 04H         | MEM_LOAD_RETIREDD.L2_MISS       | Retired loads that miss the L2 cache (precise event).       | <p>This event counts the number of retired load operations that missed the L2 cache.</p> <p>This event counts loads from cacheable memory only. It does not count loads by software prefetches.</p> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p> <p>Use IA32_PMC0 only.</p>  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                    | Definition  | Description and Comment   |
|-----------|-------------|-------------------------------|---|---|
| CBH       | 08H         | MEM_LOAD_RETIRED.L2_LINE_MISS | L2 cache line missed by retired loads (precise event).            | <p>This event counts the number of load operations that miss the L2 cache and result in a bus request to fetch the missing cache line. That is the missing cache line fetching has not yet started.</p> <p>This event count is equal to the number of cache lines fetched from memory by retired loads.</p> <p>This event counts loads from cacheable memory only. The event does not count loads by software prefetches.</p> <p>The event might not be counted if the load is blocked (see LOAD_BLOCK events).</p> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p> <p>Use IA32_PMC0 only.</p> |
| CBH       | 10H         | MEM_LOAD_RETIRED.DTLB_MISS    | Retired loads that miss the DTLB (precise event).                 | <p>This event counts the number of retired loads that missed the DTLB. The DTLB miss is not counted if the load operation causes a fault.</p> <p>This event counts loads from cacheable memory only. The event does not count loads by software prefetches.</p> <p>When this event is captured with the precise event mechanism, the collected samples contain the address of the instruction that was executed immediately after the instruction that caused the event.</p> <p>Use IA32_PMC0 only.</p>   |
| CCH       | 01H         | FP_MMX_TRANS_TO_MMX           | Transitions from Floating Point to MMX Instructions.              | This event counts the first MMX instructions following a floating-point instruction. Use this event to estimate the penalties for the transitions between floating-point and MMX states.  |
| CCH       | 02H         | FP_MMX_TRANS_TO_FP            | Transitions from MMX Instructions to Floating Point Instructions. | This event counts the first floating-point instructions following any MMX instruction. Use this event to estimate the penalties for the transitions between floating-point and MMX states.  |
| CDH       | 00H         | SIMD_ASSIST                   | SIMD assists invoked.   | This event counts the number of SIMD assists invoked. SIMD assists are invoked when an EMMS instruction is executed, changing the MMX state in the floating point stack.  |
| CEH       | 00H         | SIMD_INSTR_RETIRED            | SIMD Instructions retired.  | This event counts the number of retired SIMD instructions that use MMX registers.   |
| CFH       | 00H         | SIMD_SAT_INSTR_RETIRED        | Saturated arithmetic instructions retired.                        | This event counts the number of saturated arithmetic SIMD instructions that retired.  |
| D2H       | 01H         | RAT_STALLS.ROB_READ_PORT      | ROB read port stalls cycles.                                      | <p>This event counts the number of cycles when ROB read port stalls occurred, which did not allow new micro-ops to enter the out-of-order pipeline.</p> <p>Note that, at this stage in the pipeline, additional stalls may occur at the same cycle and prevent the stalled micro-ops from entering the pipe. In such a case, micro-ops retry entering the execution pipe in the next cycle and the ROB-read-port stall is counted again.</p>  |

Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Event Num | Umask Value | Event Name                | Definition                              | Description and Comment   |
|-----------|-------------|---------------------------|---|---|
| D2H       | 02H         | RAT_STALLS.PARTIAL_CYCLES | Partial register stall cycles.          | This event counts the number of cycles instruction execution latency became longer than the defined latency because the instruction uses a register that was partially written by previous instructions.  |
| D2H       | 04H         | RAT_STALLS.FLAGS          | Flag stall cycles.                      | This event counts the number of cycles during which execution stalled due to several reasons, one of which is a partial flag register stall.<br>A partial register stall may occur when two conditions are met: <ul style="list-style-type: none"> <li>An instruction modifies some, but not all, of the flags in the flag register.</li> <li>The next instruction, which depends on flags, depends on flags that were not modified by this instruction.</li> </ul> |
| D2H       | 08H         | RAT_STALLS.FPSW           | FPU status word stall.                  | This event indicates that the FPU status word (FPSW) is written. To obtain the number of times the FPSW is written divide the event count by 2.<br>The FPSW is written by instructions with long latency; a small count may indicate a high penalty.  |
| D2H       | 0FH         | RAT_STALLS.ANY            | All RAT stall cycles.                   | This event counts the number of stall cycles due to conditions described by: <ul style="list-style-type: none"> <li>RAT_STALLS.ROB_READ_PORT</li> <li>RAT_STALLS.PARTIAL</li> <li>RAT_STALLS.FLAGS</li> <li>RAT_STALLS.FPSW.</li> </ul>   |
| D4H       | 01H         | SEG_RENAME_STALLS.ES      | Segment rename stalls - ES.             | This event counts the number of stalls due to the lack of renaming resources for the ES segment register. If a segment is renamed, but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.   |
| D4H       | 02H         | SEG_RENAME_STALLS.DS      | Segment rename stalls - DS.             | This event counts the number of stalls due to the lack of renaming resources for the DS segment register. If a segment is renamed, but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.   |
| D4H       | 04H         | SEG_RENAME_STALLS.FS      | Segment rename stalls - FS.             | This event counts the number of stalls due to the lack of renaming resources for the FS segment register.<br>If a segment is renamed, but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.  |
| D4H       | 08H         | SEG_RENAME_STALLS.GS      | Segment rename stalls - GS.             | This event counts the number of stalls due to the lack of renaming resources for the GS segment register.<br>If a segment is renamed, but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.  |
| D4H       | 0FH         | SEG_RENAME_STALLS.ANY     | Any (ES/DS/FS/GS) segment rename stall. | This event counts the number of stalls due to the lack of renaming resources for the ES, DS, FS, and GS segment registers.<br>If a segment is renamed but not retired and a second update to the same segment occurs, a stall occurs in the front end of the pipeline until the renamed segment retires.  |



**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name                    | Definition   | Description and Comment   |
|-----------|-------------|-------------------------------|--|---|
| D5H       | 01H         | SEG_REG_RENAMES.ES            | Segment renames - ES.  | This event counts the number of times the ES segment register is renamed.   |
| D5H       | 02H         | SEG_REG_RENAMES.DS            | Segment renames - DS.  | This event counts the number of times the DS segment register is renamed.   |
| D5H       | 04H         | SEG_REG_RENAMES.FS            | Segment renames - FS.  | This event counts the number of times the FS segment register is renamed.   |
| D5H       | 08H         | SEG_REG_RENAMES.GS            | Segment renames - GS.  | This event counts the number of times the GS segment register is renamed.   |
| D5H       | 0FH         | SEG_REG_RENAMES.ANY           | Any (ES/DS/FS/GS) segment rename.  | This event counts the number of times any of the four segment registers (ES/DS/FS/GS) is renamed.   |
| DCH       | 01H         | RESOURCE_STALLS.ROB_FULL      | Cycles during which the ROB full.  | This event counts the number of cycles when the number of instructions in the pipeline waiting for retirement reaches the limit the processor can handle.<br><br>A high count for this event indicates that there are long latency operations in the pipe (possibly load and store operations that miss the L2 cache, and other instructions that depend on these cannot execute until the former instructions complete execution). In this situation new instructions cannot enter the pipe and start execution.   |
| DCH       | 02H         | RESOURCE_STALLS.RS_FULL       | Cycles during which the RS full.   | This event counts the number of cycles when the number of instructions in the pipeline waiting for execution reaches the limit the processor can handle.<br><br>A high count of this event indicates that there are long latency operations in the pipe (possibly load and store operations that miss the L2 cache, and other instructions that depend on these cannot execute until the former instructions complete execution). In this situation new instructions cannot enter the pipe and start execution.   |
| DCH       | 04          | RESOURCE_STALLS.LD_ST         | Cycles during which the pipeline has exceeded load or store limit or waiting to commit all stores. | This event counts the number of cycles while resource-related stalls occur due to: <ul style="list-style-type: none"> <li>▪ The number of load instructions in the pipeline reached the limit the processor can handle. The stall ends when a loading instruction retires.</li> <li>▪ The number of store instructions in the pipeline reached the limit the processor can handle. The stall ends when a storing instruction commits its data to the cache or memory.</li> <li>▪ There is an instruction in the pipe that can be executed only when all previous stores complete and their data is committed in the caches or memory. For example, the SFENCE and MFENCE instructions require this behavior.</li> </ul> |
| DCH       | 08H         | RESOURCE_STALLS.FPCW          | Cycles stalled due to FPU control word write.  | This event counts the number of cycles while execution was stalled due to writing the floating-point unit (FPU) control word.   |
| DCH       | 10H         | RESOURCE_STALLS.BR_MISS_CLEAR | Cycles stalled due to branch misprediction.  | This event counts the number of cycles after a branch misprediction is detected at execution until the branch and all older micro-ops retire. During this time new micro-ops cannot enter the out-of-order pipeline.  |

**Table 19-25. Performance Events in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Event Num | Umask Value | Event Name          | Definition                           | Description and Comment   |
|-----------|-------------|---------------------|--------------------------------------|---|
| DCH       | 1FH         | RESOURCE_STALLS.ANY | Resource related stalls.             | This event counts the number of cycles while resource-related stalls occurs for any conditions described by the following events: <ul style="list-style-type: none"> <li>▪ RESOURCE_STALLS.ROB_FULL</li> <li>▪ RESOURCE_STALLS.RS_FULL</li> <li>▪ RESOURCE_STALLS.LD_ST</li> <li>▪ RESOURCE_STALLS.FPCW</li> <li>▪ RESOURCE_STALLS.BR_MISS_CLEAR</li> </ul>   |
| E0H       | 00H         | BR_INST_DECODED     | Branch instructions decoded.         | This event counts the number of branch instructions decoded.  |
| E4H       | 00H         | BOGUS_BR            | Bogus branches.                      | This event counts the number of byte sequences that were mistakenly detected as taken branch instructions. This results in a BACLEAR event. This occurs mainly after task switches.   |
| E6H       | 00H         | BACLEARS            | BACLEARS asserted.                   | This event counts the number of times the front end is resteeered, mainly when the BPU cannot provide a correct prediction and this is corrected by other branch handling mechanisms at the front and. This can occur if the code has many branches such that they cannot be consumed by the BPU.<br>Each BACLEAR asserted costs approximately 7 cycles of instruction fetch. The effect on total execution time depends on the surrounding code. |
| F0H       | 00H         | PREF_RQSTS_UP       | Upward prefetches issued from DPL.   | This event counts the number of upward prefetches issued from the Data Prefetch Logic (DPL) to the L2 cache. A prefetch request issued to the L2 cache cannot be cancelled and the requested cache line is fetched to the L2 cache.   |
| F8H       | 00H         | PREF_RQSTS_DN       | Downward prefetches issued from DPL. | This event counts the number of downward prefetches issued from the Data Prefetch Logic (DPL) to the L2 cache. A prefetch request issued to the L2 cache cannot be cancelled and the requested cache line is fetched to the L2 cache.   |

## 19.13 PERFORMANCE MONITORING EVENTS FOR PROCESSORS BASED ON THE GOLDMONT PLUS MICROARCHITECTURE

Intel Atom processors based on the Goldmont Plus microarchitecture support the architectural performance monitoring events listed in Table 19-1 and fixed-function performance events using a fixed counter. They also support the following performance monitoring events listed in Table 19-27. These events apply to processors with CPUID signature of 06\_7AH. In addition, processors based on the Goldmont Plus microarchitecture also support the events listed in Table 19-27 (see Section 19.14, "Performance Monitoring Events for Processors Based on the Goldmont Microarchitecture"). For an event listed in Table 19-27 that also appears in the model-specific tables of prior generations, Table 19-27 supersedes prior generation tables.

Performance monitoring event descriptions may refer to terminology described in Section B.2, "Intel® Xeon® processor 5500 Series," in Appendix B of the *Intel® 64 and IA-32 Architectures Optimization Reference Manual*.

In Goldmont Plus microarchitecture, performance monitoring events that support Processor Event Based Sampling (PEBS) and PEBS records that contain processor state information that are associated with at-retirement tagging are marked by "Precise Event".

**Table 19-26. Performance Events for the Goldmont Plus Microarchitecture**

| Event Num. | Umask Value | Event Name                             | Description  | Comment                                      |
|------------|-------------|--|--|--|
| 00H        | 01H         | INST_RETIRED.ANY                       | Counts the number of instructions that retire execution. For instructions that consist of multiple uops, this event counts the retirement of the last uop of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers. This event uses fixed counter 0. You cannot collect a PEBS record for this event. | Fixed Event, Precise Event, Not Reduced Skid |
| 08H        | 02H         | DTLB_LOAD_MISSES.WALK_COMPLETED_4K     | Counts page walks completed due to demand data loads (including SW prefetches) whose address translations missed in all TLB levels and were mapped to 4K pages. The page walks can end with or without a page fault.   |  |
| 08H        | 04H         | DTLB_LOAD_MISSES.WALK_COMPLETED_2M_4M  | Counts page walks completed due to demand data loads (including SW prefetches) whose address translations missed in all TLB levels and were mapped to 2M or 4M pages. The page walks can end with or without a page fault.   |  |
| 08H        | 08H         | DTLB_LOAD_MISSES.WALK_COMPLETED_1GB    | Counts page walks completed due to demand data loads (including SW prefetches) whose address translations missed in all TLB levels and were mapped to 1GB pages. The page walks can end with or without a page fault.  |  |
| 08H        | 10H         | DTLB_LOAD_MISSES.WALK_PENDING          | Counts once per cycle for each page walk occurring due to a load (demand data loads or Sw prefetches). Includes cycles spent traversing the Extended Page Table (EPT). Average cycles per walk can be calculated by dividing by the number of walks.   |  |
| 49H        | 02H         | DTLB_STORE_MISSES.WALK_COMPLETED_4K    | Counts page walks completed due to demand data stores whose address translations missed in the TLB and were mapped to 4K pages. The page walks can end with or without a page fault.   |  |
| 49H        | 04H         | DTLB_STORE_MISSES.WALK_COMPLETED_2M_4M | Counts page walks completed due to demand data stores whose address translations missed in the TLB and were mapped to 2M or 4M pages. The page walks can end with or without a page fault.   |  |
| 49H        | 08H         | DTLB_STORE_MISSES.WALK_COMPLETED_1GB   | Counts page walks completed due to demand data stores whose address translations missed in the TLB and were mapped to 1GB pages. The page walks can end with or without a page fault.  |  |
| 49H        | 10H         | DTLB_STORE_MISSES.WALK_PENDING         | Counts once per cycle for each page walk occurring due to a demand data store. Includes cycles spent traversing the Extended Page Table (EPT). Average cycles per walk can be calculated by dividing by the number of walks.   |  |

**Table 19-26. Performance Events for the Goldmont Plus Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Name                       | Description  | Comment |
|------------|-------------|----------------------------------|--|---------|
| 4FH        | 10H         | EPT.WALK_PENDING                 | Counts once per cycle for each page walk only while traversing the Extended Page Table (EPT), and does not count during the rest of the translation. The EPT is used for translating Guest-Physical Addresses to Physical Addresses for Virtual Machine Monitors (VMMs). Average cycles per walk can be calculated by dividing the count by number of walks. |         |
| 85H        | 02H         | ITLB_MISSES.WALK_COMPLETED_4K    | Counts page walks completed due to instruction fetches whose address translations missed in the TLB and were mapped to 4K pages. The page walks can end with or without a page fault.  |         |
| 85H        | 04H         | ITLB_MISSES.WALK_COMPLETED_2M_4M | Counts page walks completed due to instruction fetches whose address translations missed in the TLB and were mapped to 2M or 4M pages. The page walks can end with or without a page fault.  |         |
| 85H        | 08H         | ITLB_MISSES.WALK_COMPLETED_1GB   | Counts page walks completed due to instruction fetches whose address translations missed in the TLB and were mapped to 1GB pages. The page walks can end with or without a page fault.   |         |
| 85H        | 10H         | ITLB_MISSES.WALK_PENDING         | Counts once per cycle for each page walk occurring due to an instruction fetch. Includes cycles spent traversing the Extended Page Table (EPT). Average cycles per walk can be calculated by dividing by the number of walks.  |         |
| BDH        | 20H         | TLB_FLUSHES.STLB_ANY             | Counts STLB flushes. The TLBs are flushed on instructions like INVLPG and MOV to CR3.  |         |
| C3H        | 20H         | MACHINE_CLEARS.PAGE_FAULT        | Counts the number of times that the machines clears due to a page fault. Covers both I-side and D-side (Loads/Stores) page faults. A page fault occurs when either page is not present, or an access violation.  |         |

## 19.14 PERFORMANCE MONITORING EVENTS FOR PROCESSORS BASED ON THE GOLDMONT MICROARCHITECTURE

Intel Atom processors based on the Goldmont microarchitecture support the architectural performance monitoring events listed in Table 19-1 and fixed-function performance events using a fixed counter. In addition, they also support the following model-specific performance monitoring events listed in Table 19-27. These events apply to processors with CUID signatures of 06\_5CH, 06\_5FH, and 06\_7AH.

Performance monitoring event descriptions may refer to terminology described in Section B.2, “Intel® Xeon® processor 5500 Series,” in Appendix B of the *Intel® 64 and IA-32 Architectures Optimization Reference Manual*.

In Goldmont microarchitecture, performance monitoring events that support Processor Event Based Sampling (PEBS) and PEBS records that contain processor state information that are associated with at-retirement tagging are marked by “Precise Event”.

**Table 19-27. Performance Events for the Goldmont Microarchitecture**

| Event Num. | Umask Value | Event Name              | Description  | Comment       |
|------------|-------------|-------------------------|--|---------------|
| 03H        | 10H         | LD_BLOCKS.ALL_BLOCK     | Counts anytime a load that retires is blocked for any reason.  | Precise Event |
| 03H        | 08H         | LD_BLOCKS.UTLB_MISS     | Counts loads blocked because they are unable to find their physical address in the micro TLB (UTLB).   | Precise Event |
| 03H        | 02H         | LD_BLOCKS.STORE_FORWARD | Counts a load blocked from using a store forward because of an address/size mismatch; only one of the loads blocked from each store will be counted. | Precise Event |

**Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Name                        | Description   | Comment       |
|------------|-------------|-----------------------------------|---|---------------|
| 03H        | 01H         | LD_BLOCKS.DATA_UNK NOWN           | Counts a load blocked from using a store forward, but did not occur because the store data was not available at the right time. The forward might occur subsequently when the data is available.  | Precise Event |
| 03H        | 04H         | LD_BLOCKS.4K_ALIAS                | Counts loads that block because their address modulo 4K matches a pending store.  | Precise Event |
| 05H        | 01H         | PAGE_WALKS.D_SIDE_CYCLES          | Counts every core cycle when a Data-side (walks due to data operation) page walk is in progress.  |               |
| 05H        | 02H         | PAGE_WALKS.I_SIDE_CYCLES          | Counts every core cycle when an Instruction-side (walks due to an instruction fetch) page walk is in progress.  |               |
| 05H        | 03H         | PAGE_WALKS.CYCLES                 | Counts every core cycle a page-walk is in progress due to either a data memory operation, or an instruction fetch.  |               |
| 0EH        | 00H         | UOPS_ISSUED.ANY                   | Counts uops issued by the front end and allocated into the back end of the machine. This event counts uops that retire as well as uops that were speculatively executed but didn't retire. The sort of speculative uops that might be counted includes, but is not limited to those uops issued in the shadow of a mispredicted branch, those uops that are inserted during an assist (such as for a denormal floating-point result), and (previously allocated) uops that might be canceled during a machine clear.    |               |
| 13H        | 02H         | MISALIGN_MEM_REF.LOAD_PAGE_SPLIT  | Counts when a memory load of a uop that spans a page boundary (a split) is retired.   | Precise Event |
| 13H        | 04H         | MISALIGN_MEM_REF.STORE_PAGE_SPLIT | Counts when a memory store of a uop that spans a page boundary (a split) is retired.  | Precise Event |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE       | Counts memory requests originating from the core that reference a cache line in the L2 cache.   |               |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS            | Counts memory requests originating from the core that miss in the L2 cache.   |               |
| 30H        | 00H         | L2_REJECT_XQ.ALL                  | Counts the number of demand and prefetch transactions that the L2 XQ rejects due to a full or near full condition which likely indicates back pressure from the intra-die interconnect (IDI) fabric. The XQ may reject transactions from the L2Q (non-cacheable requests), L2 misses and L2 write-back victims.   |               |
| 31H        | 00H         | CORE_REJECT_L2Q.ALL               | Counts the number of demand and L1 prefetcher requests rejected by the L2Q due to a full or nearly full condition which likely indicates back pressure from L2Q. It also counts requests that would have gone directly to the XQ, but are rejected due to a full or nearly full condition, indicating back pressure from the IDI link. The L2Q may also reject transactions from a core to ensure fairness between cores, or to delay a core's dirty eviction when the address conflicts with incoming external snoops. |               |
| 3CH        | 00H         | CPU_CLK_UNHALTED.CORE_P           | Core cycles when core is not halted. This event uses a programmable general purpose performance counter.  |               |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF              | Reference cycles when core is not halted. This event uses a programmable general purpose performance counter.   |               |
| 51H        | 01H         | DL1.DIRTY_EVICTION                | Counts when a modified (dirty) cache line is evicted from the data L1 cache and needs to be written back to memory. No count will occur if the evicted line is clean, and hence does not require a writeback.   |               |

Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)

| Event Num. | Umask Value | Event Name                             | Description   | Comment |
|------------|-------------|--|---|---------|
| 80H        | 01H         | ICACHE.HIT                             | Counts requests to the Instruction Cache (ICache) for one or more bytes in an ICache Line and that cache line is in the ICache (hit). The event strives to count on a cache line basis, so that multiple accesses which hit in a single cache line count as one ICACHE.HIT. Specifically, the event counts when straight line code crosses the cache line boundary, or when a branch target is to a new line, and that cache line is in the ICache. This event counts differently than Intel processors based on the Silvermont microarchitecture.            |         |
| 80H        | 02H         | ICACHE.MISSES                          | Counts requests to the Instruction Cache (ICache) for one or more bytes in an ICache Line and that cache line is not in the ICache (miss). The event strives to count on a cache line basis, so that multiple accesses which miss in a single cache line count as one ICACHE.MISS. Specifically, the event counts when straight line code crosses the cache line boundary, or when a branch target is to a new line, and that cache line is not in the ICache. This event counts differently than Intel processors based on the Silvermont microarchitecture. |         |
| 80H        | 03H         | ICACHE.ACCESSSES                       | Counts requests to the Instruction Cache (ICache) for one or more bytes in an ICache Line. The event strives to count on a cache line basis, so that multiple fetches to a single cache line count as one ICACHE.ACCESS. Specifically, the event counts when accesses from straight line code crosses the cache line boundary, or when a branch target is to a new line. This event counts differently than Intel processors based on the Silvermont microarchitecture.   |         |
| 81H        | 04H         | ITLB.MISS                              | Counts the number of times the machine was unable to find a translation in the Instruction Translation Lookaside Buffer (ITLB) for a linear address of an instruction fetch. It counts when new translations are filled into the ITLB. The event is speculative in nature, but will not count translations (page walks) that are begun and not finished, or translations that are finished but not filled into the ITLB.  |         |
| 86H        | 00H         | FETCH_STALL.ALL                        | Counts cycles that fetch is stalled due to any reason. That is, the decoder queue is able to accept bytes, but the fetch unit is unable to provide bytes. This will include cycles due to an ITLB miss, ICache miss and other events.   |         |
| 86H        | 01H         | FETCH_STALL.ITLB_FILL_PENDING_CYCLES   | Counts cycles that fetch is stalled due to an outstanding ITLB miss. That is, the decoder queue is able to accept bytes, but the fetch unit is unable to provide bytes due to an ITLB miss. Note: this event is not the same as page walk cycles to retrieve an instruction translation.  |         |
| 86H        | 02H         | FETCH_STALL.ICACHE_FILL_PENDING_CYCLES | Counts cycles that an ICache miss is outstanding, and instruction fetch is stalled. That is, the decoder queue is able to accept bytes, but the fetch unit is unable to provide bytes, while an ICache miss is outstanding. Note this event is not the same as cycles to retrieve an instruction due to an ICache miss. Rather, it is the part of the Instruction Cache (ICache) miss time where no bytes are available for the decoder.  |         |

**Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Name             | Description  | Comment                         |
|------------|-------------|------------------------|--|---------------------------------|
| 9CH        | 00H         | UOPS_NOT_DELIVERED.ANY | <p>This event is used to measure front-end inefficiencies, i.e., when the front end of the machine is not delivering uops to the back end and the back end has not stalled. This event can be used to identify if the machine is truly front-end bound. When this event occurs, it is an indication that the front end of the machine is operating at less than its theoretical peak performance.</p> <p>Background: We can think of the processor pipeline as being divided into 2 broader parts: the front end and the back end. The front end is responsible for fetching the instruction, decoding into uops in machine understandable format and putting them into a uop queue to be consumed by the back end. The back end then takes these uops and allocates the required resources. When all resources are ready, uops are executed. If the back end is not ready to accept uops from the front end, then we do not want to count these as front-end bottlenecks. However, whenever we have bottlenecks in the back end, we will have allocation unit stalls and eventually force the front end to wait until the back end is ready to receive more uops. This event counts only when the back end is requesting more micro-uops and the front end is not able to provide them. When 3 uops are requested and no uops are delivered, the event counts 3. When 3 are requested, and only 1 is delivered, the event counts 2. When only 2 are delivered, the event counts 1. Alternatively stated, the event will not count if 3 uops are delivered, or if the back end is stalled and not requesting any uops at all. Counts indicate missed opportunities for the front end to deliver a uop to the back end. Some examples of conditions that cause front-end inefficiencies are: lcache misses, ITLB misses, and decoder restrictions that limit the front-end bandwidth.</p> <p>Known Issues: Some uops require multiple allocation slots. These uops will not be charged as a front end 'not delivered' opportunity, and will be regarded as a back-end problem. For example, the INC instruction has one uop that requires 2 issue slots. A stream of INC instructions will not count as UOPS_NOT_DELIVERED, even though only one instruction can be issued per clock. The low uop issue rate for a stream of INC instructions is considered to be a back-end issue.</p> |                                 |
| B7H        | 01H, 02H    | OFFCORE_RESPONSE       | Requires MSR_OFFCORE_RESP[0,1] to specify request type and response. (Duplicated for both MSRs.)   |                                 |
| COH        | 00H         | INST_RETIRED.ANY_P     | <p>Counts the number of instructions that retire execution. For instructions that consist of multiple uops, this event counts the retirement of the last uop of the instruction. The event continues counting during hardware interrupts, traps, and inside interrupt handlers. This is an architectural performance event. This event uses a programmable general purpose performance counter. *This event is a Precise Event: the EventingRIP field in the PEBS record is precise to the address of the instruction which caused the event.</p> <p>Note: Because PEBS records can be collected only on IA32_PMC0, only one event can use the PEBS facility at a time.</p>  | Precise Event                   |
| C2H        | 00H         | UOPS_RETIRED.ANY       | Counts uops which have retired.  | Precise Event, Not Reduced Skid |

Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)

| Event Num. | Umask Value | Event Name                         | Description   | Comment                         |
|------------|-------------|------------------------------------|---|---------------------------------|
| C2H        | 01H         | UOPS_RETIRED.MS                    | Counts uops retired that are from the complex flows issued by the micro-sequencer (MS). Counts both the uops from a micro-coded instruction, and the uops that might be generated from a micro-coded assist.  | Precise Event, Not Reduced Skid |
| C2H        | 08H         | UOPS_RETIRED.FPDIV                 | Counts the number of floating point divide uops retired.  | Precise Event                   |
| C2H        | 10H         | UOPS_RETIRED.IDIV                  | Counts the number of integer divide uops retired.   | Precise Event                   |
| C3H        | 01H         | MACHINE_CLEARS.SMC                 | Counts the number of times that the processor detects that a program is writing to a code section and has to perform a machine clear because of that modification. Self-modifying code (SMC) causes a severe penalty in all Intel architecture processors.                          |                                 |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING     | Counts machine clears due to memory ordering issues. This occurs when a snoop request happens and the machine is uncertain if memory ordering will be preserved as another core is in the process of modifying the data.  |                                 |
| C3H        | 04H         | MACHINE_CLEARS.FP_ASSIST           | Counts machine clears due to floating-point (FP) operations needing assists. For instance, if the result was a floating-point denormal, the hardware clears the pipeline and reissues uops to produce the correct IEEE compliant denormal result.                                   |                                 |
| C3H        | 08H         | MACHINE_CLEARS.DISAMBIGUATION      | Counts machine clears due to memory disambiguation. Memory disambiguation happens when a load which has been issued conflicts with a previous un-retired store in the pipeline whose address was not known at issue time, but is later resolved to be the same as the load address. |                                 |
| C3H        | 00H         | MACHINE_CLEARS.ALL                 | Counts machine clears for any reason.   |                                 |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES       | Counts branch instructions retired for all branch types. This is an architectural performance event.  | Precise Event                   |
| C4H        | 7EH         | BR_INST_RETIRED.JCC                | Counts retired Jcc (Jump on Conditional Code/Jump if Condition is Met) branch instructions retired, including both when the branch was taken and when it was not taken.   | Precise Event                   |
| C4H        | 80H         | BR_INST_RETIRED.ALL_TAKEN_BRANCHES | Counts the number of taken branch instructions retired.   | Precise Event                   |
| C4H        | FEH         | BR_INST_RETIRED.TAKEN_JCC          | Counts Jcc (Jump on Conditional Code/Jump if Condition is Met) branch instructions retired that were taken and does not count when the Jcc branch instruction were not taken.   | Precise Event                   |
| C4H        | F9H         | BR_INST_RETIRED.CALL               | Counts near CALL branch instructions retired.   | Precise Event                   |
| C4H        | FDH         | BR_INST_RETIRED.REL_CALL           | Counts near relative CALL branch instructions retired.  | Precise Event                   |
| C4H        | FBH         | BR_INST_RETIRED.IND_CALL           | Counts near indirect CALL branch instructions retired.  | Precise Event                   |
| C4H        | F7H         | BR_INST_RETIRED.RETURN             | Counts near return branch instructions retired.   | Precise Event                   |
| C4H        | EBH         | BR_INST_RETIRED.NON_RETURN_IND     | Counts near indirect call or near indirect jmp branch instructions retired.   | Precise Event                   |
| C4H        | BFH         | BR_INST_RETIRED.FAR_BRANCH         | Counts far branch instructions retired. This includes far jump, far call and return, and Interrupt call and return.   | Precise Event                   |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES       | Counts mispredicted branch instructions retired including all branch types.   | Precise Event                   |



**Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Name                             | Description  | Comment       |
|------------|-------------|--|--|---------------|
| C5H        | 7EH         | BR_MISP_RETIRED.JCC                    | Counts mispredicted retired Jcc (Jump on Conditional Code/Jump if Condition is Met) branch instructions retired, including both when the branch was supposed to be taken and when it was not supposed to be taken (but the processor predicted the opposite condition).  | Precise Event |
| C5H        | FEH         | BR_MISP_RETIRED.TAKEN_JCC              | Counts mispredicted retired Jcc (Jump on Conditional Code/Jump if Condition is Met) branch instructions retired that were supposed to be taken but the processor predicted that it would not be taken.   | Precise Event |
| C5H        | FBH         | BR_MISP_RETIRED.IND_CALL               | Counts mispredicted near indirect CALL branch instructions retired, where the target address taken was not what the processor predicted.   | Precise Event |
| C5H        | F7H         | BR_MISP_RETIRED.RETURN                 | Counts mispredicted near RET branch instructions retired, where the return address taken was not what the processor predicted.   | Precise Event |
| C5H        | EBH         | BR_MISP_RETIRED.NON_RETURN_IND         | Counts mispredicted branch instructions retired that were near indirect call or near indirect jmp, where the target address taken was not what the processor predicted.  | Precise Event |
| CAH        | 01H         | ISSUE_SLOTS_NOT_CONSUMED.RESOURCE_FULL | Counts the number of issue slots per core cycle that were not consumed because of a full resource in the back end. Including but not limited to resources include the Re-order Buffer (ROB), reservation stations (RS), load/store buffers, physical registers, or any other needed machine resource that is currently unavailable. Note that uops must be available for consumption in order for this event to fire. If a uop is not available (Instruction Queue is empty), this event will not count. |               |
| CAH        | 02H         | ISSUE_SLOTS_NOT_CONSUMED.RECOVERY      | Counts the number of issue slots per core cycle that were not consumed by the back end because allocation is stalled waiting for a mispredicted jump to retire or other branch-like conditions (e.g. the event is relevant during certain microcode flows). Counts all issue slots blocked while within this window, including slots where uops were not available in the Instruction Queue.   |               |
| CAH        | 00H         | ISSUE_SLOTS_NOT_CONSUMED.ANY           | Counts the number of issue slots per core cycle that were not consumed by the back end due to either a full resource in the back end (RESOURCE_FULL), or due to the processor recovering from some event (RECOVERY).   |               |
| CBH        | 01H         | HW_INTERRUPTS.RECEIVED                 | Counts hardware interrupts received by the processor.  |               |
| CBH        | 02H         | HW_INTERRUPTS.MASKED                   | Counts the number of core cycles during which interrupts are masked (disabled). Increments by 1 each core cycle that EFLAGS.IF is 0, regardless of whether interrupts are pending or not.  |               |
| CBH        | 04H         | HW_INTERRUPTS.PENDING_AND_MASKED       | Counts core cycles during which there are pending interrupts, but interrupts are masked (EFLAGS.IF = 0).   |               |
| CDH        | 00H         | CYCLES_DIV_BUSY.ALL                    | Counts core cycles if either divide unit is busy.  |               |
| CDH        | 01H         | CYCLES_DIV_BUSY.IDIV                   | Counts core cycles if the integer divide unit is busy.   |               |
| CDH        | 02H         | CYCLES_DIV_BUSY.FPDIV                  | Counts core cycles if the floating point divide unit is busy.  |               |
| DOH        | 81H         | MEM_UOPS_RETIRED.ALL_LOADS             | Counts the number of load uops retired.  | Precise Event |
| DOH        | 82H         | MEM_UOPS_RETIRED.ALL_STORES            | Counts the number of store uops retired.   | Precise Event |

Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)

| Event Num. | Umask Value | Event Name                        | Description   | Comment       |
|------------|-------------|-----------------------------------|---|---------------|
| D0H        | 83H         | MEM_UOPS_RETIRED.ALL              | Counts the number of memory uops retired that are either a load or a store or both.   | Precise Event |
| D0H        | 11H         | MEM_UOPS_RETIRED.DTLB_MISS_LOADS  | Counts load uops retired that caused a DTLB miss.   | Precise Event |
| D0H        | 12H         | MEM_UOPS_RETIRED.DTLB_MISS_STORES | Counts store uops retired that caused a DTLB miss.  | Precise Event |
| D0H        | 13H         | MEM_UOPS_RETIRED.DTLB_MISS        | Counts uops retired that had a DTLB miss on load, store or either.<br>Note that when two distinct memory operations to the same page miss the DTLB, only one of them will be recorded as a DTLB miss.   | Precise Event |
| D0H        | 21H         | MEM_UOPS_RETIRED.LOCK_LOADS       | Counts locked memory uops retired. This includes 'regular' locks and bus locks. To specifically count bus locks only, see the offcore response event. A locked access is one with a lock prefix, or an exchange to memory.  | Precise Event |
| D0H        | 41H         | MEM_UOPS_RETIRED.SPLIT_LOADS      | Counts load uops retired where the data requested spans a 64 byte cache line boundary.  | Precise Event |
| D0H        | 42H         | MEM_UOPS_RETIRED.SPLIT_STORES     | Counts store uops retired where the data requested spans a 64 byte cache line boundary.   | Precise Event |
| D0H        | 43H         | MEM_UOPS_RETIRED.SPLIT            | Counts memory uops retired where the data requested spans a 64 byte cache line boundary.  | Precise Event |
| D1H        | 01H         | MEM_LOAD_UOPS_RETIRED.L1_HIT      | Counts load uops retired that hit the L1 data cache.  | Precise Event |
| D1H        | 08H         | MEM_LOAD_UOPS_RETIRED.L1_MISS     | Counts load uops retired that miss the L1 data cache.   | Precise Event |
| D1H        | 02H         | MEM_LOAD_UOPS_RETIRED.L2_HIT      | Counts load uops retired that hit in the L2 cache.  | Precise Event |
| 0xD1H      | 10H         | MEM_LOAD_UOPS_RETIRED.L2_MISS     | Counts load uops retired that miss in the L2 cache.   | Precise Event |
| D1H        | 20H         | MEM_LOAD_UOPS_RETIRED.HITM        | Counts load uops retired where the cache line containing the data was in the modified state of another core or modules cache (HITM). More specifically, this means that when the load address was checked by other caching agents (typically another processor) in the system, one of those caching agents indicated that they had a dirty copy of the data. Loads that obtain a HITM response incur greater latency than most that is typical for a load. In addition, since HITM indicates that some other processor had this data in its cache, it implies that the data was shared between processors, or potentially was a lock or semaphore value. This event is useful for locating sharing, false sharing, and contended locks. | Precise Event |

**Table 19-27. Performance Events for the Goldmont Microarchitecture (Contd.)**

| Event Num. | Umask Value | Event Name                         | Description  | Comment       |
|------------|-------------|------------------------------------|--|---------------|
| D1H        | 40H         | MEM_LOAD_UOPS_RETIRED.WCB_HIT      | Counts memory load uops retired where the data is retrieved from the WCB (or fill buffer), indicating that the load found its data while that data was in the process of being brought into the L1 cache. Typically a load will receive this indication when some other load or prefetch missed the L1 cache and was in the process of retrieving the cache line containing the data, but that process had not yet finished (and written the data back to the cache). For example, consider load X and Y, both referencing the same cache line that is not in the L1 cache. If load X misses cache first, it obtains and WCB (or fill buffer) begins the process of requesting the data. When load Y requests the data, it will either hit the WCB, or the L1 cache, depending on exactly what time the request to Y occurs. | Precise Event |
| D1H        | 80H         | MEM_LOAD_UOPS_RETIRED.DRAM_HIT     | Counts memory load uops retired where the data is retrieved from DRAM. Event is counted at retirement, so the speculative loads are ignored. A memory load can hit (or miss) the L1 cache, hit (or miss) the L2 cache, hit DRAM, hit in the WCB or receive a HITM response.  | Precise Event |
| E6H        | 01H         | BACLEARS.ALL                       | Counts the number of times a BACLEAR is signaled for any reason, including, but not limited to indirect branch/call, Jcc (Jump on Conditional Code/Jump if Condition is Met) branch, unconditional branch/call, and returns.   |               |
| E6H        | 08H         | BACLEARS.RETURN                    | Counts BACLEARS on return instructions.  |               |
| E6H        | 10H         | BACLEARS.COND                      | Counts BACLEARS on Jcc (Jump on Conditional Code/Jump if Condition is Met) branches.   |               |
| E7H        | 01H         | MS_DECODED.MS_ENTRY                | Counts the number of times the Microcode Sequencer (MS) starts a flow of uops from the MSROM. It does not count every time a uop is read from the MSROM. The most common case that this counts is when a micro-coded instruction is encountered by the front end of the machine. Other cases include when an instruction encounters a fault, trap, or microcode assist of any sort that initiates a flow of uops. The event will count MS startups for uops that are speculative, and subsequently cleared by branch mispredict or a machine clear.  |               |
| E9H        | 01H         | DECODE_RESTRICTION.PREDECODE_WRONG | Counts the number of times the prediction (from the pre-decode cache) for instruction length is incorrect.   |               |

## 19.15 PERFORMANCE MONITORING EVENTS FOR PROCESSORS BASED ON THE SILVERMONT MICROARCHITECTURE

Processors based on the Silvermont microarchitecture support the architectural performance monitoring events listed in Table 19-1 and fixed-function performance events using fixed counter. In addition, they also support the following model-specific performance monitoring events listed in Table 19-28. These processors have the CPUID signatures of 06\_37H, 06\_4AH, 06\_4DH, 06\_5AH, and 06\_5DH.

Performance monitoring event descriptions may refer to terminology described in Section B.2, "Intel® Xeon® processor 5500 Series," in Appendix B of the *Intel® 64 and IA-32 Architectures Optimization Reference Manual*.

Table 19-28. Performance Events for Silvermont Microarchitecture

| Event Num. | Umask Value | Event Name                      | Definition                                      | Description and Comment  |
|------------|-------------|---------------------------------|---|--|
| 03H        | 01H         | REHABQ.LD_BLOCK_ST_FORWARD      | Loads blocked due to store forward restriction. | This event counts the number of retired loads that were prohibited from receiving forwarded data from the store because of address mismatch.   |
| 03H        | 02H         | REHABQ.LD_BLOCK_ST_NOTREADY     | Loads blocked due to store data not ready.      | This event counts the cases where a forward was technically possible, but did not occur because the store data was not available at the right time.  |
| 03H        | 04H         | REHABQ.ST_SPLITS                | Store uops that split cache line boundary.      | This event counts the number of retire stores that experienced cache line boundary splits.   |
| 03H        | 08H         | REHABQ.LD_SPLITS                | Load uops that split cache line boundary.       | This event counts the number of retire loads that experienced cache line boundary splits.  |
| 03H        | 10H         | REHABQ.LOCK                     | Uops with lock semantics.                       | This event counts the number of retired memory operations with lock semantics. These are either implicit locked instructions such as the XCHG instruction or instructions with an explicit LOCK prefix (FOH).  |
| 03H        | 20H         | REHABQ.STA_FULL                 | Store address buffer full.                      | This event counts the number of retired stores that are delayed because there is not a store address buffer available.   |
| 03H        | 40H         | REHABQ.ANY_LD                   | Any reissued load uops.                         | This event counts the number of load uops reissued from Rehabq.  |
| 03H        | 80H         | REHABQ.ANY_ST                   | Any reissued store uops.                        | This event counts the number of store uops reissued from Rehabq.   |
| 04H        | 01H         | MEM_UOPS_RETIREDL1_MISS_LOADS   | Loads retired that missed L1 data cache.        | This event counts the number of load ops retired that miss in L1 Data cache. Note that prefetch misses will not be counted.  |
| 04H        | 02H         | MEM_UOPS_RETIREDL2_HIT_LOADS    | Loads retired that hit L2.                      | This event counts the number of load micro-ops retired that hit L2.  |
| 04H        | 04H         | MEM_UOPS_RETIREDL2_MISS_LOADS   | Loads retired that missed L2.                   | This event counts the number of load micro-ops retired that missed L2.   |
| 04H        | 08H         | MEM_UOPS_RETIREDDTLB_MISS_LOADS | Loads missed DTLB.                              | This event counts the number of load ops retired that had DTLB miss.   |
| 04H        | 10H         | MEM_UOPS_RETIREDDTLB_MISS_LOADS | Loads missed UTLB.                              | This event counts the number of load ops retired that had UTLB miss.   |
| 04H        | 20H         | MEM_UOPS_RETIREDDTLB_MISS_LOADS | Cross core or cross module hitm.                | This event counts the number of load ops retired that got data from the other core or from the other module.   |
| 04H        | 40H         | MEM_UOPS_RETIREDDTLB_MISS_LOADS | All Loads.                                      | This event counts the number of load ops retired.  |
| 04H        | 80H         | MEM_UOP_RETIREDDTLB_MISS_LOADS  | All Stores.                                     | This event counts the number of store ops retired.   |
| 05H        | 01H         | PAGE_WALKS.D_SIDE_CYCLES        | Duration of D-side page-walks in core cycles.   | This event counts every cycle when a D-side (walks due to a load) page walk is in progress. Page walk duration divided by number of page walks is the average duration of page-walks. Edge trigger bit must be cleared. Set Edge to count the number of page walks.                |
| 05H        | 02H         | PAGE_WALKS.I_SIDE_CYCLES        | Duration of I-side page-walks in core cycles.   | This event counts every cycle when an I-side (walks due to an instruction fetch) page walk is in progress. Page walk duration divided by number of page walks is the average duration of page-walks. Edge trigger bit must be cleared. Set Edge to count the number of page walks. |

**Table 19-28. Performance Events for Silvermont Microarchitecture**

| Event Num. | Umask Value | Event Name                  | Definition  | Description and Comment  |
|------------|-------------|-----------------------------|---|--|
| 05H        | 03H         | PAGE_WALKS.WALKS            | Total number of page-walks that are completed (I-side and D-side).                        | This event counts when a data (D) page walk or an instruction (I) page walk is completed or started. Since a page walk implies a TLB miss, the number of TLB misses can be counted by counting the number of pagewalks.<br>Edge trigger bit must be set. Clear Edge to count the number of cycles.   |
| 2EH        | 41H         | LONGEST_LAT_CACHE.MISS      | L2 cache request misses.  | This event counts the total number of L2 cache references and the number of L2 cache misses respectively.<br>L3 is not supported in Silvermont microarchitecture.  |
| 2EH        | 4FH         | LONGEST_LAT_CACHE.REFERENCE | L2 cache requests from this core.   | This event counts requests originating from the core that references a cache line in the L2 cache.<br>L3 is not supported in Silvermont microarchitecture.   |
| 30H        | 00H         | L2_REJECT_XQ.ALL            | Counts the number of request from the L2 that were not accepted into the XQ.              | This event counts the number of demand and prefetch transactions that the L2 XQ rejects due to a full or near full condition which likely indicates back pressure from the IDI link. The XQ may reject transactions from the L2Q (non-cacheable requests), BBS (L2 misses) and WOB (L2 write-back victims).  |
| 31H        | 00H         | CORE_REJECT_L2Q.ALL         | Counts the number of request that were not accepted into the L2Q because the L2Q is FULL. | This event counts the number of demand and L1 prefetcher requests rejected by the L2Q due to a full or nearly full condition which likely indicates back pressure from L2Q. It also counts requests that would have gone directly to the XQ, but are rejected due to a full or nearly full condition, indicating back pressure from the IDI link. The L2Q may also reject transactions from a core to insure fairness between cores, or to delay a core's dirty eviction when the address conflicts incoming external snoops. (Note that L2 prefetcher requests that are dropped are not counted by this event.) |
| 3CH        | 00H         | CPU_CLK_UNHALTED.CORE_P     | Core cycles when core is not halted.  | This event counts the number of core cycles while the core is not in a halt state. The core enters the halt state when it is running the HLT instruction. In mobile systems the core frequency may change from time to time. For this reason this event may have a changing ratio with regards to time.  |
| N/A        | N/A         | CPU_CLK_UNHALTED.CORE       | Core cycles when core is not halted.  | This uses the fixed counter 1 to count the same condition as CPU_CLK_UNHALTED.CORE_P does.   |
| 3CH        | 01H         | CPU_CLK_UNHALTED.REF_P      | Bus cycles when core is not halted.   | This event counts the number of bus cycles that the core is not in a halt state. The core enters the halt state when it is running the HLT instruction.<br>In mobile systems the core frequency may change from time. This event is not affected by core frequency changes.  |
| N/A        | N/A         | CPU_CLK_UNHALTED.REF_TSC    | Reference cycles when core is not halted.   | This event counts the number of reference cycles at a TSC rate that the core is not in a halt state. The core enters the halt state when it is running the HLT instruction.<br>In mobile systems the core frequency may change from time. This event is not affected by core frequency changes.  |
| 80H        | 01H         | ICACHE.HIT                  | Instruction fetches from Icache.  | This event counts all instruction fetches from the instruction cache.  |
| 80H        | 02H         | ICACHE.MISSES               | Icache miss.  | This event counts all instruction fetches that miss the Instruction cache or produce memory requests. This includes uncacheable fetches. An instruction fetch miss is counted only once and not once for every cycle it is outstanding.  |

Table 19-28. Performance Events for Silvermont Microarchitecture

| Event Num. | Umask Value | Event Name                      | Definition   | Description and Comment   |
|------------|-------------|---------------------------------|--|---|
| 80H        | 03H         | ICACHE.ACCESSES                 | Instruction fetches.                                     | This event counts all instruction fetches, including uncacheable fetches.   |
| B7H        | 01H         | OFFCORE_RESPONSE_0              | See Section 18.5.2.2.                                    | Requires MSR_OFFCORE_RESP0 to specify request type and response.  |
| B7H        | 02H         | OFFCORE_RESPONSE_1              | See Section 18.5.2.2.                                    | Requires MSR_OFFCORE_RESP1 to specify request type and response.  |
| C0H        | 00H         | INST_RETIRED.ANY_P              | Instructions retired (PEBS supported with IA32_PMC0).    | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers. |
| N/A        | N/A         | INST_RETIRED.ANY                | Instructions retired.                                    | This uses the fixed counter 0 to count the same condition as INST_RETIRED.ANY_P does.   |
| C2H        | 01H         | UOPS_RETIRED.MS                 | MSROM micro-ops retired.                                 | This event counts the number of micro-ops retired that were supplied from MSROM.  |
| C2H        | 10H         | UOPS_RETIRED.ALL                | Micro-ops retired.                                       | This event counts the number of micro-ops retired.  |
| C3H        | 01H         | MACHINE_CLEARS.SMC              | Self-Modifying Code detected.                            | This event counts the number of times that a program writes to a code section. Self-modifying code causes a severe penalty in all Intel® architecture processors.   |
| C3H        | 02H         | MACHINE_CLEARS.MEMORY_ORDERING  | Stalls due to Memory ordering.                           | This event counts the number of times that pipeline was cleared due to memory ordering issues.  |
| C3H        | 04H         | MACHINE_CLEARS.FP_ASSIST        | Stalls due to FP assists.                                | This event counts the number of times that pipeline stalled due to FP operations needing assists.   |
| C3H        | 08H         | MACHINE_CLEARS.ALL              | Stalls due to any causes.                                | This event counts the number of times that pipeline stalled due to due to any causes (including SMC, MO, FP assist, etc.).  |
| C4H        | 00H         | BR_INST_RETIRED.ALL_BRANCHES    | Retired branch instructions.                             | This event counts the number of branch instructions retired.  |
| C4H        | 7EH         | BR_INST_RETIRED.JCC             | Retired branch instructions that were conditional jumps. | This event counts the number of branch instructions retired that were conditional jumps.  |
| C4H        | BFH         | BR_INST_RETIRED.FAR_BRANCH      | Retired far branch instructions.                         | This event counts the number of far branch instructions retired.  |
| C4H        | EBH         | BR_INST_RETIRED.NO_N_RETURN_IND | Retired instructions of near indirect jmp or call.       | This event counts the number of branch instructions retired that were near indirect call or near indirect jmp.  |
| C4H        | F7H         | BR_INST_RETIRED.RETURN          | Retired near return instructions.                        | This event counts the number of near RET branch instructions retired.   |
| C4H        | F9H         | BR_INST_RETIRED.CALL            | Retired near call instructions.                          | This event counts the number of near CALL branch instructions retired.  |
| C4H        | FBH         | BR_INST_RETIRED.IND_CALL        | Retired near indirect call instructions.                 | This event counts the number of near indirect CALL branch instructions retired.   |
| C4H        | FDH         | BR_INST_RETIRED.REL_CALL        | Retired near relative call instructions.                 | This event counts the number of near relative CALL branch instructions retired.   |
| C4H        | FEH         | BR_INST_RETIRED.TAKEN_JCC       | Retired conditional jumps that were taken.               | This event counts the number of branch instructions retired that were conditional jumps and taken.  |
| C5H        | 00H         | BR_MISP_RETIRED.ALL_BRANCHES    | Retired mispredicted branch instructions.                | This event counts the number of mispredicted branch instructions retired.   |

**Table 19-28. Performance Events for Silvermont Microarchitecture**

| Event Num. | Umask Value | Event Name                     | Definition  | Description and Comment  |
|------------|-------------|--------------------------------|---|--|
| C5H        | 7EH         | BR_MISP_RETIRED.JCC            | Retired mispredicted conditional jumps.   | This event counts the number of mispredicted branch instructions retired that were conditional jumps.  |
| C5H        | BFH         | BR_MISP_RETIRED.FAR            | Retired mispredicted far branch instructions.   | This event counts the number of mispredicted far branch instructions retired.  |
| C5H        | EBH         | BR_MISP_RETIRED.NON_RETURN_IND | Retired mispredicted instructions of near indirect jmp or call.   | This event counts the number of mispredicted branch instructions retired that were near indirect call or near indirect jmp.                  |
| C5H        | F7H         | BR_MISP_RETIRED.RETURN         | Retired mispredicted near return instructions.  | This event counts the number of mispredicted near RET branch instructions retired.   |
| C5H        | F9H         | BR_MISP_RETIRED.CALL           | Retired mispredicted near call instructions.  | This event counts the number of mispredicted near CALL branch instructions retired.  |
| C5H        | FBH         | BR_MISP_RETIRED.IND_CALL       | Retired mispredicted near indirect call instructions.   | This event counts the number of mispredicted near indirect CALL branch instructions retired.   |
| C5H        | FDH         | BR_MISP_RETIRED.REL_CALL       | Retired mispredicted near relative call instructions  | This event counts the number of mispredicted near relative CALL branch instructions retired.   |
| C5H        | FEH         | BR_MISP_RETIRED.TAKEN_JCC      | Retired mispredicted conditional jumps that were taken.   | This event counts the number of mispredicted branch instructions retired that were conditional jumps and taken.                              |
| CAH        | 01H         | NO_ALLOC_CYCLES.ROB_FULL       | Counts the number of cycles when no uops are allocated and the ROB is full (less than 2 entries available). | Counts the number of cycles when no uops are allocated and the ROB is full (less than 2 entries available).                                  |
| CAH        | 20H         | NO_ALLOC_CYCLES.RAT_STALL      | Counts the number of cycles when no uops are allocated and a RATstall is asserted.                          | Counts the number of cycles when no uops are allocated and a RATstall is asserted.   |
| CAH        | 3FH         | NO_ALLOC_CYCLES.AL             | Front end not delivering.   | This event counts the number of cycles when the front end does not provide any instructions to be allocated for any reason.                  |
| CAH        | 50H         | NO_ALLOC_CYCLES.NOT_DELIVERED  | Front end not delivering back end not stalled.  | This event counts the number of cycles when the front end does not provide any instructions to be allocated but the back end is not stalled. |
| CBH        | 01H         | RS_FULL_STALL.MEC              | MEC RS full.  | This event counts the number of cycles the allocation pipe line stalled due to the RS for the MEC cluster is full.                           |
| CBH        | 1FH         | RS_FULL_STALL.ALL              | Any RS full.  | This event counts the number of cycles that the allocation pipe line stalled due to any one of the RS is full.                               |
| CDH        | 01H         | CYCLES_DIV_BUSY.ANY            | Divider Busy.   | This event counts the number of cycles the divider is busy.  |
| E6H        | 01H         | BACLEARS.ALL                   | BACLEARS asserted for any branch.   | This event counts the number of baclears for any type of branch.   |
| E6H        | 08H         | BACLEARS.RETURN                | BACLEARS asserted for return branch.  | This event counts the number of baclears for return branches.  |

**Table 19-28. Performance Events for Silvermont Microarchitecture**

| Event Num. | Umask Value | Event Name          | Definition                                | Description and Comment  |
|------------|-------------|---------------------|---|--|
| E6H        | 10H         | BACLEARS.COND       | BACLEARS asserted for conditional branch. | This event counts the number of baclears for conditional branches.     |
| E7H        | 01H         | MS_DECODED.MS_ENTRY | MS Decode starts.                         | This event counts the number of times the MSROM starts a flow of UOPS. |

### 19.15.1 Performance Monitoring Events for Processors Based on the Airmont Microarchitecture

Intel processors based on the Airmont microarchitecture support the same architectural and the model-specific performance monitoring events as processors based on the Silvermont microarchitecture. All of the events listed in Table 19-28 apply. These processors have the CPUID signatures that include 06\_4CH.

## 19.16 PERFORMANCE MONITORING EVENTS FOR 45 NM AND 32 NM INTEL® ATOM™ PROCESSORS

45 nm and 32 nm processors based on the Intel® Atom™ microarchitecture support the architectural performance monitoring events listed in Table 19-1 and fixed-function performance events using fixed counter listed in Table 19-24. In addition, they also support the following model-specific performance monitoring events listed in Table 19-29.

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors**

| Event Num. | Umask Value | Event Name                | Definition   | Description and Comment   |
|------------|-------------|---------------------------|--|---|
| 02H        | 81H         | STORE_FORWARDS.GO<br>OD   | Good store forwards.   | This event counts the number of times store data was forwarded directly to a load.  |
| 06H        | 00H         | SEGMENT_REG_<br>LOADS.ANY | Number of segment register loads.  | This event counts the number of segment register load operations. Instructions that load new values into segment registers cause a penalty. This event indicates performance issues in 16-bit code. If this event occurs frequently, it may be useful to calculate the number of instructions retired per segment register load. If the resulting calculation is low (on average a small number of instructions are executed between segment register loads), then the code's segment register usage should be optimized.<br><br>As a result of branch misprediction, this event is speculative and may include segment register loads that do not actually occur. However, most segment register loads are internally serialized and such speculative effects are minimized. |
| 07H        | 01H         | PREFETCH.PREFETCH<br>T0   | Streaming SIMD Extensions (SSE) PrefetchT0 instructions executed.                | This event counts the number of times the SSE instruction prefetchT0 is executed. This instruction prefetches the data to the L1 data cache and L2 cache.   |
| 07H        | 06H         | PREFETCH.SW_L2            | Streaming SIMD Extensions (SSE) PrefetchT1 and PrefetchT2 instructions executed. | This event counts the number of times the SSE instructions prefetchT1 and prefetchT2 are executed. These instructions prefetch the data to the L2 cache.  |



**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value | Event Name                      | Definition  | Description and Comment   |
|------------|-------------|---------------------------------|---|---|
| 07H        | 08H         | PREFETCH.PREFETCHNTA            | Streaming SIMD Extensions (SSE) Prefetch NTA instructions executed. | This event counts the number of times the SSE instruction prefetchNTA is executed. This instruction prefetches the data to the L1 data cache.   |
| 08H        | 07H         | DATA_TLB_MISSES.DTLB_MISS       | Memory accesses that missed the DTLB.                               | This event counts the number of Data Table Lookaside Buffer (DTLB) misses. The count includes misses detected as a result of speculative accesses. Typically a high count for this event indicates that the code accesses a large number of data pages.   |
| 08H        | 05H         | DATA_TLB_MISSES.DTLB_MISS_LD    | DTLB misses due to load operations.                                 | This event counts the number of Data Table Lookaside Buffer (DTLB) misses due to load operations. This count includes misses detected as a result of speculative accesses.  |
| 08H        | 09H         | DATA_TLB_MISSES.LO_DTLB_MISS_LD | LO_DTLB misses due to load operations.                              | This event counts the number of LO_DTLB misses due to load operations. This count includes misses detected as a result of speculative accesses.   |
| 08H        | 06H         | DATA_TLB_MISSES.DTLB_MISS_ST    | DTLB misses due to store operations.                                | This event counts the number of Data Table Lookaside Buffer (DTLB) misses due to store operations. This count includes misses detected as a result of speculative accesses.   |
| 0CH        | 03H         | PAGE_WALKS.WALKS                | Number of page-walks executed.                                      | This event counts the number of page-walks executed due to either a DTLB or ITLB miss. The page walk duration, PAGE_WALKS.CYCLES, divided by number of page walks is the average duration of a page walk. This can hint to whether most of the page-walks are satisfied by the caches or cause an L2 cache miss.<br>Edge trigger bit must be set.                       |
| 0CH        | 03H         | PAGE_WALKS.CYCLES               | Duration of page-walks in core cycles.                              | This event counts the duration of page-walks in core cycles. The paging mode in use typically affects the duration of page walks. Page walk duration divided by number of page walks is the average duration of page-walks. This can hint at whether most of the page-walks are satisfied by the caches or cause an L2 cache miss.<br>Edge trigger bit must be cleared. |
| 10H        | 01H         | X87_COMP_OPS_EXE.ANY.S          | Floating point computational micro-ops executed.                    | This event counts the number of x87 floating point computational micro-ops executed.  |
| 10H        | 81H         | X87_COMP_OPS_EXE.ANY.AR         | Floating point computational micro-ops retired.                     | This event counts the number of x87 floating point computational micro-ops retired.   |
| 11H        | 01H         | FP_ASSIST                       | Floating point assists.   | This event counts the number of floating point operations executed that required micro-code assist intervention. These assists are required in the following cases.<br>X87 instructions:<br>1. NaN or denormal are loaded to a register or used as input from memory.<br>2. Division by 0.<br>3. Underflow output.  |

Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)

| Event Num. | Umask Value                     | Event Name      | Definition                            | Description and Comment  |
|------------|---------------------------------|-----------------|---------------------------------------|--|
| 11H        | 81H                             | FP_ASSIST.AR    | Floating point assists.               | This event counts the number of floating point operations executed that required micro-code assist intervention. These assists are required in the following cases.<br>X87 instructions:<br>1. NaN or denormal are loaded to a register or used as input from memory.<br>2. Division by 0.<br>3. Underflow output.   |
| 12H        | 01H                             | MUL.S           | Multiply operations executed.         | This event counts the number of multiply operations executed. This includes integer as well as floating point multiply operations.   |
| 12H        | 81H                             | MUL.AR          | Multiply operations retired.          | This event counts the number of multiply operations retired. This includes integer as well as floating point multiply operations.  |
| 13H        | 01H                             | DIV.S           | Divide operations executed.           | This event counts the number of divide operations executed. This includes integer divides, floating point divides and square-root operations executed.   |
| 13H        | 81H                             | DIV.AR          | Divide operations retired.            | This event counts the number of divide operations retired. This includes integer divides, floating point divides and square-root operations executed.  |
| 14H        | 01H                             | CYCLES_DIV_BUSY | Cycles the divider is busy.           | This event counts the number of cycles the divider is busy executing divide or square root operations. The divide can be integer, X87 or Streaming SIMD Extensions (SSE). The square root operation can be either X87 or SSE.  |
| 21H        | See Table 18-61                 | L2_ADS          | Cycles L2 address bus is in use.      | This event counts the number of cycles the L2 address bus is being used for accesses to the L2 cache or bus queue.<br>This event can count occurrences for this core or both cores.  |
| 22H        | See Table 18-61                 | L2_DBUS_BUSY    | Cycles the L2 cache data bus is busy. | This event counts core cycles during which the L2 cache data bus is busy transferring data from the L2 cache to the core. It counts for all L1 cache misses (data and instruction) that hit the L2 cache. The count will increment by two for a full cache-line request.   |
| 24H        | See Table 18-61 and Table 18-63 | L2_LINES_IN     | L2 cache misses.                      | This event counts the number of cache lines allocated in the L2 cache. Cache lines are allocated in the L2 cache as a result of requests from the L1 data and instruction caches and the L2 hardware prefetchers to cache lines that are missing in the L2 cache.<br>This event can count occurrences for this core or both cores. This event can also count demand requests and L2 hardware prefetch requests together or separately. |
| 25H        | See Table 18-61                 | L2_M_LINES_IN   | L2 cache line modifications.          | This event counts whenever a modified cache line is written back from the L1 data cache to the L2 cache.<br>This event can count occurrences for this core or both cores.  |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value                                  | Event Name     | Definition                                | Description and Comment  |
|------------|--|----------------|---|--|
| 26H        | See Table 18-61 and Table 18-63              | L2_LINES_OUT   | L2 cache lines evicted.                   | This event counts the number of L2 cache lines evicted.<br>This event can count occurrences for this core or both cores. This event can also count evictions due to demand requests and L2 hardware prefetch requests together or separately.  |
| 27H        | See Table 18-61 and Table 18-63              | L2_M_LINES_OUT | Modified lines evicted from the L2 cache. | This event counts the number of L2 modified cache lines evicted. These lines are written back to memory unless they also exist in a shared-state in one of the L1 data caches.<br>This event can count occurrences for this core or both cores. This event can also count evictions due to demand requests and L2 hardware prefetch requests together or separately.   |
| 28H        | See Table 18-61 and Table 18-64              | L2_IFETCH      | L2 cacheable instruction fetch requests.  | This event counts the number of instruction cache line requests from the ICache. It does not include fetch requests from uncacheable memory. It does not include ITLB miss accesses.<br>This event can count occurrences for this core or both cores. This event can also count accesses to cache lines at different MESI states.  |
| 29H        | See Table 18-61, Table 18-63 and Table 18-64 | L2_LD          | L2 cache reads.                           | This event counts L2 cache read requests coming from the L1 data cache and L2 prefetchers.<br>This event can count occurrences for this core or both cores. This event can count occurrences<br>- for this core or both cores.<br>- due to demand requests and L2 hardware prefetch requests together or separately.<br>- of accesses to cache lines at different MESI states.   |
| 2AH        | See Table 18-61 and Table 18-64              | L2_ST          | L2 store requests.                        | This event counts all store operations that miss the L1 data cache and request the data from the L2 cache.<br>This event can count occurrences for this core or both cores. This event can also count accesses to cache lines at different MESI states.  |
| 2BH        | See Table 18-61 and Table 18-64              | L2_LOCK        | L2 locked accesses.                       | This event counts all locked accesses to cache lines that miss the L1 data cache.<br>This event can count occurrences for this core or both cores. This event can also count accesses to cache lines at different MESI states.   |
| 2EH        | See Table 18-61, Table 18-63 and Table 18-64 | L2_RQSTS       | L2 cache requests.                        | This event counts all completed L2 cache requests. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, instruction fetches, and all L2 hardware prefetch requests.<br>This event can count occurrences<br>- for this core or both cores.<br>- due to demand requests and L2 hardware prefetch requests together, or separately.<br>- of accesses to cache lines at different MESI states. |

Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)

| Event Num. | Umask Value                                  | Event Name                   | Definition   | Description and Comment  |
|------------|--|------------------------------|--|--|
| 2EH        | 41H  | L2_RQSTS.SELF.DEMAND.I_STATE | L2 cache demand requests from this core that missed the L2.          | This event counts all completed L2 cache demand requests from this core that miss the L2 cache. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, and instruction fetches.<br>This is an architectural performance event.   |
| 2EH        | 4FH  | L2_RQSTS.SELF.DEMAND.MESI    | L2 cache demand requests from this core.                             | This event counts all completed L2 cache demand requests from this core. This includes L1 data cache reads, writes, and locked accesses, L1 data prefetch requests, and instruction fetches.<br>This is an architectural performance event.  |
| 30H        | See Table 18-61, Table 18-63 and Table 18-64 | L2_REJECT_BUSQ               | Rejected L2 cache requests.  | This event indicates that a pending L2 cache request that requires a bus transaction is delayed from moving to the bus queue. Some of the reasons for this event are:<br>- The bus queue is full.<br>- The bus queue already holds an entry for a cache line in the same set.<br>The number of events is greater or equal to the number of requests that were rejected.<br>- For this core or both cores.<br>- Due to demand requests and L2 hardware prefetch requests together, or separately.<br>- Of accesses to cache lines at different MESI states.   |
| 32H        | See Table 18-61                              | L2_NO_REQ                    | Cycles no L2 cache requests are pending.                             | This event counts the number of cycles that no L2 cache requests are pending.  |
| 3AH        | 00H  | EIST_TRANS                   | Number of Enhanced Intel SpeedStep(R) Technology (EIST) transitions. | This event counts the number of Enhanced Intel SpeedStep(R) Technology (EIST) transitions that include a frequency change, either with or without VID change. This event is incremented only while the counting core is in C0 state. In situations where an EIST transition was caused by hardware as a result of CxE state transitions, those EIST transitions will also be registered in this event.<br><br>Enhanced Intel Speedstep Technology transitions are commonly initiated by OS, but can be initiated by HW internally. For example: CxE states are C-states (C1,C2,C3...) which not only place the CPU into a sleep state by turning off the clock and other components, but also lower the voltage (which reduces the leakage power consumption). The same is true for thermal throttling transition which uses Enhanced Intel Speedstep Technology internally. |
| 3BH        | COH  | THERMAL_TRIP                 | Number of thermal trips.   | This event counts the number of thermal trips. A thermal trip occurs whenever the processor temperature exceeds the thermal trip threshold temperature. Following a thermal trip, the processor automatically reduces frequency and voltage. The processor checks the temperature every millisecond, and returns to normal when the temperature falls below the thermal trip threshold temperature.  |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value                      | Event Name                | Definition  | Description and Comment   |
|------------|----------------------------------|---------------------------|---|---|
| 3CH        | 00H                              | CPU_CLK_UNHALTED.CORE_P   | Core cycles when core is not halted.                    | <p>This event counts the number of core cycles while the core is not in a halt state. The core enters the halt state when it is running the HLT instruction. This event is a component in many key event ratios.</p> <p>In mobile systems the core frequency may change from time to time. For this reason this event may have a changing ratio with regards to time. In systems with a constant core frequency, this event can give you a measurement of the elapsed time while the core was not in halt state by dividing the event count by the core frequency.</p> <ul style="list-style-type: none"> <li>-This is an architectural performance event.</li> <li>- The event CPU_CLK_UNHALTED.CORE_P is counted by a programmable counter.</li> <li>- The event CPU_CLK_UNHALTED.CORE is counted by a designated fixed counter, leaving the two programmable counters available for other events.</li> </ul> |
| 3CH        | 01H                              | CPU_CLK_UNHALTED.BUS      | Bus cycles when core is not halted.                     | <p>This event counts the number of bus cycles while the core is not in the halt state. This event can give you a measurement of the elapsed time while the core was not in the halt state, by dividing the event count by the bus frequency. The core enters the halt state when it is running the HLT instruction.</p> <p>The event also has a constant ratio with CPU_CLK_UNHALTED.REF event, which is the maximum bus to processor frequency ratio.</p> <p>Non-halted bus cycles are a component in many key event ratios.</p>   |
| 3CH        | 02H                              | CPU_CLK_UNHALTED.NO_OTHER | Bus cycles when core is active and the other is halted. | <p>This event counts the number of bus cycles during which the core remains non-halted, and the other core on the processor is halted.</p> <p>This event can be used to determine the amount of parallelism exploited by an application or a system. Divide this event count by the bus frequency to determine the amount of time that only one core was in use.</p>  |
| 40H        | 21H                              | L1D_CACHE.LD              | L1 Cacheable Data Reads.                                | This event counts the number of data reads from cacheable memory.   |
| 40H        | 22H                              | L1D_CACHE.ST              | L1 Cacheable Data Writes.                               | This event counts the number of data writes to cacheable memory.  |
| 60H        | See Table 18-61 and Table 18-62. | BUS_REQUEST_OUTSTANDING   | Outstanding cacheable data read bus requests duration.  | This event counts the number of pending full cache line read transactions on the bus occurring in each cycle. A read transaction is pending from the cycle it is sent on the bus until the full cache line is received by the processor. NOTE: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.   |

Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)

| Event Num. | Umask Value                      | Event Name      | Definition                                 | Description and Comment  |
|------------|----------------------------------|-----------------|--|--|
| 61H        | See Table 18-62.                 | BUS_BNR_DRV     | Number of Bus Not Ready signals asserted.  | <p>This event counts the number of Bus Not Ready (BNR) signals that the processor asserts on the bus to suspend additional bus requests by other bus agents. A bus agent asserts the BNR signal when the number of data and snoop transactions is close to the maximum that the bus can handle.</p> <p>While this signal is asserted, new transactions cannot be submitted on the bus. As a result, transaction latency may have higher impact on program performance. NOTE: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.</p>      |
| 62H        | See Table 18-62.                 | BUS_DRDY_CLOCKS | Bus cycles when data is sent on the bus.   | <p>This event counts the number of bus cycles during which the DRDY (Data Ready) signal is asserted on the bus. The DRDY signal is asserted when data is sent on the bus.</p> <p>This event counts the number of bus cycles during which this agent (the processor) writes data on the bus back to memory or to other bus agents. This includes all explicit and implicit data writebacks, as well as partial writes.</p> <p>Note: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.</p>  |
| 63H        | See Table 18-61 and Table 18-62. | BUS_LOCK_CLOCKS | Bus cycles when a LOCK signal is asserted. | <p>This event counts the number of bus cycles, during which the LOCK signal is asserted on the bus. A LOCK signal is asserted when there is a locked memory access, due to:</p> <ul style="list-style-type: none"> <li>- Uncacheable memory.</li> <li>- Locked operation that spans two cache lines.</li> <li>- Page-walk from an uncacheable page table.</li> </ul> <p>Bus locks have a very high performance penalty and it is highly recommended to avoid such accesses. NOTE: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.</p> |
| 64H        | See Table 18-61.                 | BUS_DATA_RCV    | Bus cycles while processor receives data.  | <p>This event counts the number of cycles during which the processor is busy receiving data. NOTE: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.</p>  |
| 65H        | See Table 18-61 and Table 18-62. | BUS_TRANS_BRD   | Burst read bus transactions.               | <p>This event counts the number of burst read transactions including:</p> <ul style="list-style-type: none"> <li>- L1 data cache read misses (and L1 data cache hardware prefetches).</li> <li>- L2 hardware prefetches by the DPL and L2 streamer.</li> <li>- IFU read misses of cacheable lines.</li> </ul> <p>It does not include RFO transactions.</p>   |
| 66H        | See Table 18-61 and Table 18-62. | BUS_TRANS_RFO   | RFO bus transactions.                      | <p>This event counts the number of Read For Ownership (RFO) bus transactions, due to store operations that miss the L1 data cache and the L2 cache. This event also counts RFO bus transactions due to locked operations.</p>  |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value                      | Event Name        | Definition                                | Description and Comment   |
|------------|----------------------------------|-------------------|---|---|
| 67H        | See Table 18-61 and Table 18-62. | BUS_TRANS_WB      | Explicit writeback bus transactions.      | This event counts all explicit writeback bus transactions due to dirty line evictions. It does not count implicit writebacks due to invalidation by a snoop request.  |
| 68H        | See Table 18-61 and Table 18-62. | BUS_TRANS_IFETCH  | Instruction-fetch bus transactions.       | This event counts all instruction fetch full cache line bus transactions.   |
| 69H        | See Table 18-61 and Table 18-62. | BUS_TRANS_INVALID | Invalidate bus transactions.              | This event counts all invalidate transactions. Invalidate transactions are generated when:<br>- A store operation hits a shared line in the L2 cache.<br>- A full cache line write misses the L2 cache or hits a shared line in the L2 cache. |
| 6AH        | See Table 18-61 and Table 18-62. | BUS_TRANS_PWR     | Partial write bus transaction.            | This event counts partial write bus transactions.   |
| 6BH        | See Table 18-61 and Table 18-62. | BUS_TRANS_P       | Partial bus transactions.                 | This event counts all (read and write) partial bus transactions.  |
| 6CH        | See Table 18-61 and Table 18-62. | BUS_TRANS_IO      | IO bus transactions.                      | This event counts the number of completed I/O bus transactions as a result of IN and OUT instructions. The count does not include memory mapped IO.   |
| 6DH        | See Table 18-61 and Table 18-62. | BUS_TRANS_DEF     | Deferred bus transactions.                | This event counts the number of deferred transactions.  |
| 6EH        | See Table 18-61 and Table 18-62. | BUS_TRANS_BURST   | Burst (full cache-line) bus transactions. | This event counts burst (full cache line) transactions including:<br>- Burst reads.<br>- RFOs.<br>- Explicit writebacks.<br>- Write combine lines.  |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value                      | Event Name       | Definition                            | Description and Comment  |
|------------|----------------------------------|------------------|---------------------------------------|--|
| 6FH        | See Table 18-61 and Table 18-62. | BUS_TRANS_MEM    | Memory bus transactions.              | This event counts all memory bus transactions including:<br>- Burst transactions.<br>- Partial reads and writes.<br>- Invalidate transactions.<br>The BUS_TRANS_MEM count is the sum of BUS_TRANS_BURST, BUS_TRANS_P and BUS_TRANS_INVALID.  |
| 70H        | See Table 18-61 and Table 18-62. | BUS_TRANS_ANY    | All bus transactions.                 | This event counts all bus transactions. This includes:<br>- Memory transactions.<br>- IO transactions (non memory-mapped).<br>- Deferred transaction completion.<br>- Other less frequent transactions, such as interrupts.  |
| 77H        | See Table 18-61 and Table 18-64. | EXT_SNOOP        | External snoops.                      | This event counts the snoop responses to bus transactions. Responses can be counted separately by type and by bus agent.<br>Note: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.   |
| 7AH        | See Table 18-62.                 | BUS_HIT_DRV      | HIT signal asserted.                  | This event counts the number of bus cycles during which the processor drives the HIT# pin to signal HIT snoop response.<br>Note: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.  |
| 7BH        | See Table 18-62.                 | BUS_HITM_DRV     | HITM signal asserted.                 | This event counts the number of bus cycles during which the processor drives the HITM# pin to signal HITM snoop response.<br>NOTE: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.  |
| 7DH        | See Table 18-61.                 | BUSQ_EMPTY       | Bus queue is empty.                   | This event counts the number of cycles during which the core did not have any pending transactions in the bus queue.<br>Note: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled.   |
| 7EH        | See Table 18-61 and Table 18-62. | SNOOP_STALL_DRV  | Bus stalled for snoops.               | This event counts the number of times that the bus snoop stall signal is asserted. During the snoop stall cycles no new bus transactions requiring a snoop response can be initiated on the bus.<br>Note: This event is thread-independent and will not provide a count per logical processor when AnyThr is disabled. |
| 7FH        | See Table 18-61.                 | BUS_IO_WAIT      | IO requests waiting in the bus queue. | This event counts the number of core cycles during which IO requests wait in the bus queue. This event counts IO requests from the core.   |
| 80H        | 03H                              | ICACHE.ACCESSSES | Instruction fetches.                  | This event counts all instruction fetches, including uncacheable fetches.  |
| 80H        | 02H                              | ICACHE.MISSES    | Icache miss.                          | This event counts all instruction fetches that miss the Instruction cache or produce memory requests. This includes uncacheable fetches. An instruction fetch miss is counted only once and not once for every cycle it is outstanding.  |
| 82H        | 04H                              | ITLB.FLUSH       | ITLB flushes.                         | This event counts the number of ITLB flushes.  |
| 82H        | 02H                              | ITLB.MISSES      | ITLB misses.                          | This event counts the number of instruction fetches that miss the ITLB.  |



**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value | Event Name                       | Definition                                    | Description and Comment   |
|------------|-------------|----------------------------------|---|---|
| AAH        | 02H         | MACRO_INSTS.CISC_DECODED         | CISC macro instructions decoded.              | This event counts the number of complex instructions decoded, but not necessarily executed or retired. Only one complex instruction can be decoded at a time.   |
| AAH        | 03H         | MACRO_INSTS.ALL_DECODED          | All Instructions decoded.                     | This event counts the number of instructions decoded.   |
| B0H        | 00H         | SIMD_UOPS_EXEC.S                 | SIMD micro-ops executed (excluding stores).   | This event counts all the SIMD micro-ops executed. This event does not count MOVQ and MOVD stores from register to memory.  |
| B0H        | 80H         | SIMD_UOPS_EXEC.AR                | SIMD micro-ops retired (excluding stores).    | This event counts the number of SIMD saturated arithmetic micro-ops executed.   |
| B1H        | 00H         | SIMD_SAT_UOP_EXEC.S              | SIMD saturated arithmetic micro-ops executed. | This event counts the number of SIMD saturated arithmetic micro-ops executed.   |
| B1H        | 80H         | SIMD_SAT_UOP_EXEC.AR             | SIMD saturated arithmetic micro-ops retired.  | This event counts the number of SIMD saturated arithmetic micro-ops retired.  |
| B3H        | 01H         | SIMD_UOP_TYPE_EXEC.MUL.S         | SIMD packed multiply micro-ops executed.      | This event counts the number of SIMD packed multiply micro-ops executed.  |
| B3H        | 81H         | SIMD_UOP_TYPE_EXEC.MUL.AR        | SIMD packed multiply micro-ops retired.       | This event counts the number of SIMD packed multiply micro-ops retired.   |
| B3H        | 02H         | SIMD_UOP_TYPE_EXEC.SHIFT.S       | SIMD packed shift micro-ops executed.         | This event counts the number of SIMD packed shift micro-ops executed.   |
| B3H        | 82H         | SIMD_UOP_TYPE_EXEC.SHIFT.AR      | SIMD packed shift micro-ops retired.          | This event counts the number of SIMD packed shift micro-ops retired.  |
| B3H        | 04H         | SIMD_UOP_TYPE_EXEC.PACK.S        | SIMD pack micro-ops executed.                 | This event counts the number of SIMD pack micro-ops executed.   |
| B3H        | 84H         | SIMD_UOP_TYPE_EXEC.PACK.AR       | SIMD pack micro-ops retired.                  | This event counts the number of SIMD pack micro-ops retired.  |
| B3H        | 08H         | SIMD_UOP_TYPE_EXEC.UNPACK.S      | SIMD unpack micro-ops executed.               | This event counts the number of SIMD unpack micro-ops executed.   |
| B3H        | 88H         | SIMD_UOP_TYPE_EXEC.UNPACK.AR     | SIMD unpack micro-ops retired.                | This event counts the number of SIMD unpack micro-ops retired.  |
| B3H        | 10H         | SIMD_UOP_TYPE_EXEC.LOGICAL.S     | SIMD packed logical micro-ops executed.       | This event counts the number of SIMD packed logical micro-ops executed.   |
| B3H        | 90H         | SIMD_UOP_TYPE_EXEC.LOGICAL.AR    | SIMD packed logical micro-ops retired.        | This event counts the number of SIMD packed logical micro-ops retired.  |
| B3H        | 20H         | SIMD_UOP_TYPE_EXEC.ARITHMETIC.S  | SIMD packed arithmetic micro-ops executed.    | This event counts the number of SIMD packed arithmetic micro-ops executed.  |
| B3H        | A0H         | SIMD_UOP_TYPE_EXEC.ARITHMETIC.AR | SIMD packed arithmetic micro-ops retired.     | This event counts the number of SIMD packed arithmetic micro-ops retired.   |
| COH        | 00H         | INST_RETIRED.ANY_P               | Instructions retired (precise event).         | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers. |

Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)

| Event Num. | Umask Value | Event Name                        | Definition  | Description and Comment  |
|------------|-------------|-----------------------------------|---|--|
| N/A        | 00H         | INST_RETIRED.ANY                  | Instructions retired.   | This event counts the number of instructions that retire execution. For instructions that consist of multiple micro-ops, this event counts the retirement of the last micro-op of the instruction. The counter continues counting during hardware interrupts, traps, and inside interrupt handlers.  |
| C2H        | 10H         | UOPS_RETIRED.ANY                  | Micro-ops retired.  | This event counts the number of micro-ops retired. The processor decodes complex macro instructions into a sequence of simpler micro-ops. Most instructions are composed of one or two micro-ops. Some instructions are decoded into longer sequences such as repeat instructions, floating point transcendental instructions, and assists. In some cases micro-op sequences are fused or whole instructions are fused into one micro-op. See other UOPS_RETIRED events for differentiating retired fused and non-fused micro-ops.   |
| C3H        | 01H         | MACHINE_CLEAR.SMC                 | Self-Modifying Code detected.                                 | This event counts the number of times that a program writes to a code section. Self-modifying code causes a severe penalty in all Intel® architecture processors.  |
| C4H        | 00H         | BR_INST_RETIRED.ANY               | Retired branch instructions.                                  | This event counts the number of branch instructions retired.<br><b>This is an architectural performance event.</b>   |
| C4H        | 01H         | BR_INST_RETIRED.PRED_NOT_TAKEN    | Retired branch instructions that were predicted not-taken.    | This event counts the number of branch instructions retired that were correctly predicted to be not-taken.   |
| C4H        | 02H         | BR_INST_RETIRED.MISPRED_NOT_TAKEN | Retired branch instructions that were mispredicted not-taken. | This event counts the number of branch instructions retired that were mispredicted and not-taken.  |
| C4H        | 04H         | BR_INST_RETIRED.PRED_TAKEN        | Retired branch instructions that were predicted taken.        | This event counts the number of branch instructions retired that were correctly predicted to be taken.   |
| C4H        | 08H         | BR_INST_RETIRED.MISPRED_TAKEN     | Retired branch instructions that were mispredicted taken.     | This event counts the number of branch instructions retired that were mispredicted and taken.  |
| C4H        | 0AH         | BR_INST_RETIRED.MISPRED           | Retired mispredicted branch instructions (precise event).     | This event counts the number of retired branch instructions that were mispredicted by the processor. A branch misprediction occurs when the processor predicts that the branch would be taken, but it is not, or vice-versa. Mispredicted branches degrade the performance because the processor starts executing instructions along a wrong path it predicts. When the misprediction is discovered, all the instructions executed in the wrong path must be discarded, and the processor must start again on the correct path.<br><br>Using the Profile-Guided Optimization (PGO) features of the Intel® C++ compiler may help reduce branch mispredictions. See the compiler documentation for more information on this feature. |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value | Event Name                                      | Definition  | Description and Comment  |
|------------|-------------|---|---|--|
|            |             |   |   | <p>To determine the branch misprediction ratio, divide the BR_INST_RETIREDMISPRED event count by the number of BR_INST_RETIREDAANY event count. To determine the number of mispredicted branches per instruction, divide the number of mispredicted branches by the INST_RETIREDAANY event count. To measure the impact of the branch mispredictions use the event RESOURCE_STALLS.BR_MISS_CLEAR.</p> <p>Tips:</p> <ul style="list-style-type: none"> <li>- See the optimization guide for tips on reducing branch mispredictions.</li> <li>- PGO's purpose is to have straight line code for the most frequent execution paths, reducing branches taken and increasing the "basic block" size, possibly also reducing the code footprint or working-set.</li> </ul>   |
| C4H        | 0CH         | BR_INST_RETIREDTAKEN                            | Retired taken branch instructions.                        | This event counts the number of branches retired that were taken.  |
| C4H        | 0FH         | BR_INST_RETIREDAANY1                            | Retired branch instructions.                              | This event counts the number of branch instructions retired that were mispredicted. This event is a duplicate of BR_INST_RETIREDMISPRED.   |
| C5H        | 00H         | BR_INST_RETIREDMISPRED                          | Retired mispredicted branch instructions (precise event). | <p>This event counts the number of retired branch instructions that were mispredicted by the processor. A branch misprediction occurs when the processor predicts that the branch would be taken, but it is not, or vice-versa. Mispredicted branches degrade the performance because the processor starts executing instructions along a wrong path it predicts. When the misprediction is discovered, all the instructions executed in the wrong path must be discarded, and the processor must start again on the correct path.</p> <p>Using the Profile-Guided Optimization (PGO) features of the Intel® C++ compiler may help reduce branch mispredictions. See the compiler documentation for more information on this feature.</p> <p>To determine the branch misprediction ratio, divide the BR_INST_RETIREDMISPRED event count by the number of BR_INST_RETIREDAANY event count. To determine the number of mispredicted branches per instruction, divide the number of mispredicted branches by the INST_RETIREDAANY event count. To measure the impact of the branch mispredictions use the event RESOURCE_STALLS.BR_MISS_CLEAR.</p> <p>Tips:</p> <ul style="list-style-type: none"> <li>- See the optimization guide for tips on reducing branch mispredictions.</li> <li>- PGO's purpose is to have straight line code for the most frequent execution paths, reducing branches taken and increasing the "basic block" size, possibly also reducing the code footprint or working-set.</li> </ul> |
| C6H        | 01H         | CYCLES_INT_MASKED.CYCLES_INT_MASKED             | Cycles during which interrupts are disabled.              | This event counts the number of cycles during which interrupts are disabled.   |
| C6H        | 02H         | CYCLES_INT_MASKED.CYCLES_INT_PENDING_AND_MASKED | Cycles during which interrupts are pending and disabled.  | This event counts the number of cycles during which there are pending interrupts but interrupts are disabled.  |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value | Event Name                           | Definition   | Description and Comment  |
|------------|-------------|--------------------------------------|--|--|
| C7H        | 01H         | SIMD_INST_RETIRED.PACKED_SINGLE      | Retired Streaming SIMD Extensions (SSE) packed-single instructions.                  | This event counts the number of SSE packed-single instructions retired.  |
| C7H        | 02H         | SIMD_INST_RETIRED.SCALAR_SINGLE      | Retired Streaming SIMD Extensions (SSE) scalar-single instructions.                  | This event counts the number of SSE scalar-single instructions retired.  |
| C7H        | 04H         | SIMD_INST_RETIRED.PACKED_DOUBLE      | Retired Streaming SIMD Extensions 2 (SSE2) packed-double instructions.               | This event counts the number of SSE2 packed-double instructions retired.   |
| C7H        | 08H         | SIMD_INST_RETIRED.SCALAR_DOUBLE      | Retired Streaming SIMD Extensions 2 (SSE2) scalar-double instructions.               | This event counts the number of SSE2 scalar-double instructions retired.   |
| C7H        | 10H         | SIMD_INST_RETIRED.VECTOR             | Retired Streaming SIMD Extensions 2 (SSE2) vector instructions.                      | This event counts the number of SSE2 vector instructions retired.  |
| C7H        | 1FH         | SIMD_INST_RETIRED.ANY                | Retired Streaming SIMD instructions.   | This event counts the overall number of SIMD instructions retired. To count each type of SIMD instruction separately, use the following events:<br>SIMD_INST_RETIRED.PACKED_SINGLE<br>SIMD_INST_RETIRED.SCALAR_SINGLE<br>SIMD_INST_RETIRED.PACKED_DOUBLE<br>SIMD_INST_RETIRED.SCALAR_DOUBLE<br>SIMD_INST_RETIRED.VECTOR. |
| C8H        | 00H         | HW_INT_RCV                           | Hardware interrupts received.  | This event counts the number of hardware interrupts received by the processor. This event will count twice for dual-pipe micro-ops.  |
| CAH        | 01H         | SIMD_COMP_INST_RETIRED.PACKED_SINGLE | Retired computational Streaming SIMD Extensions (SSE) packed-single instructions.    | This event counts the number of computational SSE packed-single instructions retired. Computational instructions perform arithmetic computations, like add, multiply and divide. Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.         |
| CAH        | 02H         | SIMD_COMP_INST_RETIRED.SCALAR_SINGLE | Retired computational Streaming SIMD Extensions (SSE) scalar-single instructions.    | This event counts the number of computational SSE scalar-single instructions retired. Computational instructions perform arithmetic computations, like add, multiply and divide. Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.         |
| CAH        | 04H         | SIMD_COMP_INST_RETIRED.PACKED_DOUBLE | Retired computational Streaming SIMD Extensions 2 (SSE2) packed-double instructions. | This event counts the number of computational SSE2 packed-double instructions retired. Computational instructions perform arithmetic computations, like add, multiply and divide. Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.        |

**Table 19-29. Performance Events for 45 nm, 32 nm Intel® Atom™ Processors (Contd.)**

| Event Num. | Umask Value | Event Name                           | Definition   | Description and Comment  |
|------------|-------------|--------------------------------------|--|--|
| CAH        | 08H         | SIMD_COMP_INST_RETIRED.SCALAR_DOUBLE | Retired computational Streaming SIMD Extensions 2 (SSE2) scalar-double instructions. | This event counts the number of computational SSE2 scalar-double instructions retired. Computational instructions perform arithmetic computations, like add, multiply and divide. Instructions that perform load and store operations or logical operations, like XOR, OR, and AND are not counted by this event.  |
| CBH        | 01H         | MEM_LOAD_RETIRED.L2_HIT              | Retired loads that hit the L2 cache (precise event).                                 | This event counts the number of retired load operations that missed the L1 data cache and hit the L2 cache.  |
| CBH        | 02H         | MEM_LOAD_RETIRED.L2_MISS             | Retired loads that miss the L2 cache (precise event).                                | This event counts the number of retired load operations that missed the L2 cache.  |
| CBH        | 04H         | MEM_LOAD_RETIRED.DTLB_MISS           | Retired loads that miss the DTLB (precise event).                                    | This event counts the number of retired loads that missed the DTLB. The DTLB miss is not counted if the load operation causes a fault.   |
| CDH        | 00H         | SIMD_ASSIST                          | SIMD assists invoked.  | This event counts the number of SIMD assists invoked. SIMD assists are invoked when an EMMS instruction is executed after MMX™ technology code has changed the MMX state in the floating point stack. For example, these assists are required in the following cases.<br>Streaming SIMD Extensions (SSE) instructions:<br>1. Denormal input when the DAZ (Denormals Are Zeros) flag is off.<br>2. Underflow result when the FTZ (Flush To Zero) flag is off. |
| CEH        | 00H         | SIMD_INSTR_RETIRED                   | SIMD Instructions retired.   | This event counts the number of SIMD instructions that retired.  |
| CFH        | 00H         | SIMD_SAT_INSTR_RETIRED               | Saturated arithmetic instructions retired.   | This event counts the number of saturated arithmetic SIMD instructions that retired.   |
| E0H        | 01H         | BR_INST_DECODED                      | Branch instructions decoded.   | This event counts the number of branch instructions decoded.   |
| E4H        | 01H         | BOGUS_BR                             | Bogus branches.  | This event counts the number of byte sequences that were mistakenly detected as taken branch instructions. This results in a BACLEAR event and the BTB is flushed. This occurs mainly after task switches.   |
| E6H        | 01H         | BACLEAR.ANY                          | BACLEARs asserted.   | This event counts the number of times the front end is redirected for a branch prediction, mainly when an early branch prediction is corrected by other branch handling mechanisms in the front end. This can occur if the code has many branches such that they cannot be consumed by the branch predictor. Each Baclear asserted costs approximately 7 cycles. The effect on total execution time depends on the surrounding code.                         |

## 19.17 PERFORMANCE MONITORING EVENTS FOR INTEL® CORE™ SOLO AND INTEL® CORE™ DUO PROCESSORS

Table 19-30 lists model-specific performance events for Intel® Core™ Duo processors. If a model-specific event requires qualification in core specificity, it is indicated in the comment column. Table 19-30 also applies to Intel® Core™ Solo processors; bits in the unit mask corresponding to core-specificity are reserved and should be 00B.

**Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors**

| Event Num. | Event Mask Mnemonic | Umask Value | Description   | Comment  |
|------------|---------------------|-------------|---|--|
| 03H        | LD_Blocks           | 00H         | Load operations delayed due to store buffer blocks. The preceding store may be blocked due to unknown address, unknown data, or conflict due to partial overlap between the load and store.           |  |
| 04H        | SD_Drains           | 00H         | Cycles while draining store buffers.  |  |
| 05H        | Misalign_Mem_Ref    | 00H         | Misaligned data memory references (MOB splits of loads and stores).   |  |
| 06H        | Seg_Reg_Loads       | 00H         | Segment register loads.   |  |
| 07H        | SSE_PrefNta_Ret     | 00H         | SSE software prefetch instruction PREFETCHNTA retired.  |  |
| 07H        | SSE_PrefT1_Ret      | 01H         | SSE software prefetch instruction PREFETCHT1 retired.   |  |
| 07H        | SSE_PrefT2_Ret      | 02H         | SSE software prefetch instruction PREFETCHT2 retired.   |  |
| 07H        | SSE_NTStores_Ret    | 03H         | SSE streaming store instruction retired.  |  |
| 10H        | FP_Comps_Op_Exec    | 00H         | FP computational instruction executed. FADD, FSUB, FCOM, FMULs, MUL, IMUL, FDIVs, DIV, IDIV, FPREMs, FSQRT are included; but exclude FADD or FMUL used in the middle of a transcendental instruction. |  |
| 11H        | FP_Assist           | 00H         | FP exceptions experienced microcode assists.  | IA32_PMC1 only.  |
| 12H        | Mul                 | 00H         | Multiply operations (a speculative count, including FP and integer multiplies).   | IA32_PMC1 only.  |
| 13H        | Div                 | 00H         | Divide operations (a speculative count, including FP and integer divisions).  | IA32_PMC1 only.  |
| 14H        | Cycles_Div_Busy     | 00H         | Cycles the divider is busy.   | IA32_PMC0 only.  |
| 21H        | L2_ADS              | 00H         | L2 Address strobos.   | Requires core-specificity.                               |
| 22H        | Dbus_Busy           | 00H         | Core cycle during which data bus was busy (increments by 4).  | Requires core-specificity.                               |
| 23H        | Dbus_Busy_Rd        | 00H         | Cycles data bus is busy transferring data to a core (increments by 4).  | Requires core-specificity.                               |
| 24H        | L2_Lines_In         | 00H         | L2 cache lines allocated.   | Requires core-specificity and HW prefetch qualification. |
| 25H        | L2_M_Lines_In       | 00H         | L2 Modified-state cache lines allocated.  | Requires core-specificity.                               |
| 26H        | L2_Lines_Out        | 00H         | L2 cache lines evicted.   | Requires core-specificity and HW prefetch qualification. |
| 27H        | L2_M_Lines_Out      | 00H         | L2 Modified-state cache lines evicted.  |  |

**Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors (Contd.)**

| Event Num. | Event Mask Mnemonic     | Umask Value                 | Description  | Comment   |
|------------|-------------------------|-----------------------------|--|---|
| 28H        | L2_IFetch               | Requires MESI qualification | L2 instruction fetches from instruction fetch unit (includes speculative fetches). | Requires core-specificity.                            |
| 29H        | L2_LD                   | Requires MESI qualification | L2 cache reads.  | Requires core-specificity.                            |
| 2AH        | L2_ST                   | Requires MESI qualification | L2 cache writes (includes speculation).  | Requires core-specificity.                            |
| 2EH        | L2_Rqsts                | Requires MESI qualification | L2 cache reference requests.   | Requires core-specificity, HW prefetch qualification. |
| 30H        | L2_Reject_Cycles        | Requires MESI qualification | Cycles L2 is busy and rejecting new requests.                                      |   |
| 32H        | L2_No_Request_Cycles    | Requires MESI qualification | Cycles there is no request to access L2.   |   |
| 3AH        | EST_Trans_All           | 00H                         | Any Intel Enhanced SpeedStep(R) Technology transitions.                            |   |
| 3AH        | EST_Trans_All           | 10H                         | Intel Enhanced SpeedStep Technology frequency transitions.                         |   |
| 3BH        | Thermal_Trip            | C0H                         | Duration in a thermal trip based on the current core clock.                        | Use edge trigger to count occurrence.                 |
| 3CH        | NonHlt_Ref_Cycles       | 01H                         | Non-halted bus cycles.   |   |
| 3CH        | Serial_Execution_Cycles | 02H                         | Non-halted bus cycles of this core executing code while the other core is halted.  |   |
| 40H        | DCache_Cache_LD         | Requires MESI qualification | L1 cacheable data read operations.   |   |
| 41H        | DCache_Cache_ST         | Requires MESI qualification | L1 cacheable data write operations.  |   |
| 42H        | DCache_Cache_Lock       | Requires MESI qualification | L1 cacheable lock read operations to invalid state.                                |   |
| 43H        | Data_Mem_Ref            | 01H                         | L1 data read and writes of cacheable and non-cacheable types.                      |   |
| 44H        | Data_Mem_Cache_Ref      | 02H                         | L1 data cacheable read and write operations.                                       |   |
| 45H        | DCache_Repl             | 0FH                         | L1 data cache line replacements.   |   |
| 46H        | DCache_M_Repl           | 00H                         | L1 data M-state cache line allocated.  |   |
| 47H        | DCache_M_Evict          | 00H                         | L1 data M-state cache line evicted.  |   |
| 48H        | DCache_Pend_Miss        | 00H                         | Weighted cycles of L1 miss outstanding.  | Use Cmask =1 to count duration.                       |
| 49H        | Dtlb_Miss               | 00H                         | Data references that missed TLB.   |   |
| 4BH        | SSE_PrefNta_Miss        | 00H                         | PREFETCHNTA missed all caches.   |   |
| 4BH        | SSE_PrefT1_Miss         | 01H                         | PREFETCHT1 missed all caches.  |   |
| 4BH        | SSE_PrefT2_Miss         | 02H                         | PREFETCHT2 missed all caches.  |   |
| 4BH        | SSE_NTStores_Miss       | 03H                         | SSE streaming store instruction missed all caches.                                 |   |
| 4FH        | L1_Pref_Req             | 00H                         | L1 prefetch requests due to DCU cache misses.                                      | May overcount if request re-submitted.                |

Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors (Contd.)

| Event Num. | Event Mask Mnemonic | Umask Value  | Description  | Comment   |
|------------|---------------------|--|--|---|
| 60H        | Bus_Req_Outstanding | 00; Requires core-specificity, and agent specificity | Weighted cycles of cacheable bus data read requests. This event counts full-line read request from DCU or HW prefetcher, but not RFO, write, instruction fetches, or others. | Use Cmask =1 to count duration.<br>Use Umask bit 12 to include HWP or exclude HWP separately.     |
| 61H        | Bus_BNR_Clocks      | 00H  | External bus cycles while BNR asserted.  |   |
| 62H        | Bus_DRDY_Clocks     | 00H  | External bus cycles while DRDY asserted.   | Requires agent specificity.   |
| 63H        | Bus_Locks_Clocks    | 00H  | External bus cycles while bus lock signal asserted.  | Requires core specificity.  |
| 64H        | Bus_Data_Rcv        | 40H  | Number of data chunks received by this processor.  |   |
| 65H        | Bus_Trans_Brd       | See comment.   | Burst read bus transactions (data or code).  | Requires core specificity.  |
| 66H        | Bus_Trans_RFO       | See comment.   | Completed read for ownership (RFO) transactions.   | Requires agent specificity.   |
| 68H        | Bus_Trans_Ifetch    | See comment.   | Completed instruction fetch transactions.  |   |
| 69H        | Bus_Trans_Inval     | See comment.   | Completed invalidate transactions.   | Requires core specificity.  |
| 6AH        | Bus_Trans_Pwr       | See comment.   | Completed partial write transactions.  | Each transaction counts its address strobe.<br>Retried transaction may be counted more than once. |
| 6BH        | Bus_Trans_P         | See comment.   | Completed partial transactions (include partial read + partial write + line write).  |   |
| 6CH        | Bus_Trans_IO        | See comment.   | Completed I/O transactions (read and write).   |   |
| 6DH        | Bus_Trans_Def       | 20H  | Completed defer transactions.  | Requires core specificity.<br>Retried transaction may be counted more than once.                  |
| 67H        | Bus_Trans_WB        | COH  | Completed writeback transactions from DCU (does not include L2 writebacks).  | Requires agent specificity.   |
| 6EH        | Bus_Trans_Burst     | COH  | Completed burst transactions (full line transactions include reads, write, RFO, and writebacks).   | Each transaction counts its address strobe.   |
| 6FH        | Bus_Trans_Mem       | COH  | Completed memory transactions. This includes Bus_Trans_Burst + Bus_Trans_P+Bus_Trans_Inval.  | Retried transaction may be counted more than once.  |
| 70H        | Bus_Trans_Any       | COH  | Any completed bus transactions.  |   |
| 77H        | Bus_Snoops          | 00H  | Counts any snoop on the bus.   | Requires MESI qualification.<br>Requires agent specificity.                                       |
| 78H        | DCU_Snoop_To_Share  | 01H  | DCU snoops to share-state L1 cache line due to L1 misses.  | Requires core specificity.  |
| 7DH        | Bus_Not_In_Use      | 00H  | Number of cycles there is no transaction from the core.  | Requires core specificity.  |
| 7EH        | Bus_Snoop_Stall     | 00H  | Number of bus cycles while bus snoop is stalled.   |   |
| 80H        | ICache_Reads        | 00H  | Number of instruction fetches from ICache, streaming buffers (both cacheable and uncacheable fetches).   |   |



**Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors (Contd.)**

| Event Num. | Event Mask Mnemonic   | Umask Value | Description   | Comment |
|------------|-----------------------|-------------|---|---------|
| 81H        | ICache_Misses         | 00H         | Number of instruction fetch misses from ICache, streaming buffers.  |         |
| 85H        | ITLB_Misses           | 00H         | Number of iTLB misses.  |         |
| 86H        | IFU_Mem_Stall         | 00H         | Cycles IFU is stalled while waiting for data from memory.   |         |
| 87H        | ILD_Stall             | 00H         | Number of instruction length decoder stalls (Counts number of LCP stalls).  |         |
| 88H        | Br_Inst_Exec          | 00H         | Branch instruction executed (includes speculation).   |         |
| 89H        | Br_Missp_Exec         | 00H         | Branch instructions executed and mispredicted at execution (includes branches that do not have prediction or mispredicted). |         |
| 8AH        | Br_BAC_Missp_Exec     | 00H         | Branch instructions executed that were mispredicted at front end.   |         |
| 8BH        | Br_Cnd_Exec           | 00H         | Conditional branch instructions executed.   |         |
| 8CH        | Br_Cnd_Missp_Exec     | 00H         | Conditional branch instructions executed that were mispredicted.  |         |
| 8DH        | Br_Ind_Exec           | 00H         | Indirect branch instructions executed.  |         |
| 8EH        | Br_Ind_Missp_Exec     | 00H         | Indirect branch instructions executed that were mispredicted.   |         |
| 8FH        | Br_Ret_Exec           | 00H         | Return branch instructions executed.  |         |
| 90H        | Br_Ret_Missp_Exec     | 00H         | Return branch instructions executed that were mispredicted.   |         |
| 91H        | Br_Ret_BAC_Missp_Exec | 00H         | Return branch instructions executed that were mispredicted at the front end.  |         |
| 92H        | Br_Call_Exec          | 00H         | Return call instructions executed.  |         |
| 93H        | Br_Call_Missp_Exec    | 00H         | Return call instructions executed that were mispredicted.   |         |
| 94H        | Br_Ind_Call_Exec      | 00H         | Indirect call branch instructions executed.   |         |
| A2H        | Resource_Stall        | 00H         | Cycles while there is a resource related stall (renaming, buffer entries) as seen by allocator.                             |         |
| B0H        | MMX_Instr_Exec        | 00H         | Number of MMX instructions executed (does not include MOVQ and MOVD stores).  |         |
| B1H        | SIMD_Int_Sat_Exec     | 00H         | Number of SIMD Integer saturating instructions executed.  |         |
| B3H        | SIMD_Int_Pmul_Exec    | 01H         | Number of SIMD Integer packed multiply instructions executed.   |         |
| B3H        | SIMD_Int_Psft_Exec    | 02H         | Number of SIMD Integer packed shift instructions executed.  |         |
| B3H        | SIMD_Int_Pck_Exec     | 04H         | Number of SIMD Integer pack operations instruction executed.  |         |
| B3H        | SIMD_Int_Upck_Exec    | 08H         | Number of SIMD Integer unpack instructions executed.  |         |
| B3H        | SIMD_Int_Plog_Exec    | 10H         | Number of SIMD Integer packed logical instructions executed.  |         |
| B3H        | SIMD_Int_Pari_Exec    | 20H         | Number of SIMD Integer packed arithmetic instructions executed.   |         |

Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors (Contd.)

| Event Num. | Event Mask Mnemonic       | Umask Value | Description  | Comment             |
|------------|---------------------------|-------------|--|---------------------|
| C0H        | Instr_Ret                 | 00H         | Number of instruction retired (Macro fused instruction count as 2).                                      |                     |
| C1H        | FP_Comp_Instr_Ret         | 00H         | Number of FP compute instructions retired (X87 instruction or instruction that contains X87 operations). | Use IA32_PMC0 only. |
| C2H        | Uops_Ret                  | 00H         | Number of micro-ops retired (include fused uops).  |                     |
| C3H        | SMC_Detected              | 00H         | Number of times self-modifying code condition detected.  |                     |
| C4H        | Br_Instr_Ret              | 00H         | Number of branch instructions retired.   |                     |
| C5H        | Br_MisPred_Ret            | 00H         | Number of mispredicted branch instructions retired.  |                     |
| C6H        | Cycles_Int_Masked         | 00H         | Cycles while interrupt is disabled.  |                     |
| C7H        | Cycles_Int_Pedning_Masked | 00H         | Cycles while interrupt is disabled and interrupts are pending.   |                     |
| C8H        | HW_Int_Rx                 | 00H         | Number of hardware interrupts received.  |                     |
| C9H        | Br_Taken_Ret              | 00H         | Number of taken branch instruction retired.  |                     |
| CAH        | Br_MisPred_Taken_Ret      | 00H         | Number of taken and mispredicted branch instructions retired.  |                     |
| CCH        | MMX_FP_Trans              | 00H         | Number of transitions from MMX to X87.   |                     |
| CCH        | FP_MMX_Trans              | 01H         | Number of transitions from X87 to MMX.   |                     |
| CDH        | MMX_Assist                | 00H         | Number of EMMS executed.   |                     |
| CEH        | MMX_Instr_Ret             | 00H         | Number of MMX instruction retired.   |                     |
| D0H        | Instr_Decoded             | 00H         | Number of instruction decoded.   |                     |
| D7H        | ESP_Uops                  | 00H         | Number of ESP folding instruction decoded.   |                     |
| D8H        | SIMD_FP_SP_Ret            | 00H         | Number of SSE/SSE2 single precision instructions retired (packed and scalar).                            |                     |
| D8H        | SIMD_FP_SP_S_Ret          | 01H         | Number of SSE/SSE2 scalar single precision instructions retired.   |                     |
| D8H        | SIMD_FP_DP_P_Ret          | 02H         | Number of SSE/SSE2 packed double precision instructions retired.   |                     |
| D8H        | SIMD_FP_DP_S_Ret          | 03H         | Number of SSE/SSE2 scalar double precision instructions retired.   |                     |
| D8H        | SIMD_Int_128_Ret          | 04H         | Number of SSE2 128 bit integer instructions retired.   |                     |
| D9H        | SIMD_FP_SP_P_Comp_Ret     | 00H         | Number of SSE/SSE2 packed single precision compute instructions retired (does not include AND, OR, XOR). |                     |
| D9H        | SIMD_FP_SP_S_Comp_Ret     | 01H         | Number of SSE/SSE2 scalar single precision compute instructions retired (does not include AND, OR, XOR). |                     |
| D9H        | SIMD_FP_DP_P_Comp_Ret     | 02H         | Number of SSE/SSE2 packed double precision compute instructions retired (does not include AND, OR, XOR). |                     |
| D9H        | SIMD_FP_DP_S_Comp_Ret     | 03H         | Number of SSE/SSE2 scalar double precision compute instructions retired (does not include AND, OR, XOR). |                     |

**Table 19-30. Performance Events in Intel® Core™ Solo and Intel® Core™ Duo Processors (Contd.)**

| Event Num. | Event Mask Mnemonic | Umask Value | Description  | Comment |
|------------|---------------------|-------------|--|---------|
| DAH        | Fused_Uops_Ret      | 00H         | All fused uops retired.  |         |
| DAH        | Fused_Ld_Uops_Ret   | 01H         | Fused load uops retired.   |         |
| DAH        | Fused_St_Uops_Ret   | 02H         | Fused store uops retired.  |         |
| DBH        | Unfusion            | 00H         | Number of unfusion events in the ROB (due to exception).         |         |
| E0H        | Br_Instr_Decoded    | 00H         | Branch instructions decoded.                                     |         |
| E2H        | BTB_Misses          | 00H         | Number of branches the BTB did not produce a prediction.         |         |
| E4H        | Br_Bogus            | 00H         | Number of bogus branches.  |         |
| E6H        | BAClears            | 00H         | Number of BAClears asserted.                                     |         |
| F0H        | Pref_Rqsts_Up       | 00H         | Number of hardware prefetch requests issued in forward streams.  |         |
| F8H        | Pref_Rqsts_Dn       | 00H         | Number of hardware prefetch requests issued in backward streams. |         |

## 19.18 PENTIUM® 4 AND INTEL® XEON® PROCESSOR PERFORMANCE MONITORING EVENTS

Tables 19-31, 19-32 and 19-33 list performance monitoring events that can be counted or sampled on processors based on Intel NetBurst® microarchitecture. Table 19-31 lists the non-retirement events, and Table 19-32 lists the at-retirement events. Tables 19-34, 19-35, and 19-36 describes three sets of parameters that are available for three of the at-retirement counting events defined in Table 19-32. Table 19-37 shows which of the non-retirement and at retirement events are logical processor specific (TS) (see Section 18.6.4.4, "Performance Monitoring Events") and which are non-logical processor specific (TI).

Some of the Pentium 4 and Intel Xeon processor performance monitoring events may be available only to specific models. The performance monitoring events listed in Tables 19-31 and 19-32 apply to processors with CPUID signature that matches family encoding 15, model encoding 0, 1, 2 3, 4, or 6. Table applies to processors with a CPUID signature that matches family encoding 15, model encoding 3, 4 or 6.

The functionality of performance monitoring events in Pentium 4 and Intel Xeon processors is also available when IA-32e mode is enabled.

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting**

| Event Name      | Event Parameters         | Parameter Value              | Description   |
|-----------------|--------------------------|------------------------------|---|
| TC_deliver_mode |                          |                              | This event counts the duration (in clock cycles) of the operating modes of the trace cache and decode engine in the processor package. The mode is specified by one or more of the event mask bits. |
|                 | ESCR restrictions        | MSR_TC_ESCR0<br>MSR_TC_ESCR1 |   |
|                 | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7   |   |
|                 | ESCR Event Select        | 01H                          | ESCR[31:25]   |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name        | Event Parameters         | Parameter Value                | Description   |
|-------------------|--------------------------|--------------------------------|---|
|                   | ESCR Event Mask          | Bit                            | ESCR[24:9]  |
|                   |                          | 0: DD                          | Both logical processors are in deliver mode.  |
|                   |                          | 1: DB                          | Logical processor 0 is in deliver mode and logical processor 1 is in build mode.  |
|                   |                          | 2: DI                          | Logical processor 0 is in deliver mode and logical processor 1 is either halted, under a machine clear condition or transitioning to a long microcode flow.   |
|                   |                          | 3: BD                          | Logical processor 0 is in build mode and logical processor 1 is in deliver mode.  |
|                   |                          | 4: BB                          | Both logical processors are in build mode.  |
|                   |                          | 5: BI                          | Logical processor 0 is in build mode and logical processor 1 is either halted, under a machine clear condition or transitioning to a long microcode flow.   |
|                   |                          | 6: ID                          | Logical processor 0 is either halted, under a machine clear condition or transitioning to a long microcode flow. Logical processor 1 is in deliver mode.  |
|                   |                          | 7: IB                          | Logical processor 0 is either halted, under a machine clear condition or transitioning to a long microcode flow. Logical processor 1 is in build mode.  |
|                   |                          | CCCR Select                    | 01H   |
|                   | Event Specific Notes     |                                | If only one logical processor is available from a physical processor package, the event mask should be interpreted as logical processor 1 is halted. Event mask bit 2 was previously known as "DELIVER", bit 5 was previously known as "BUILD". |
| BPU_fetch_request |                          |                                | This event counts instruction fetch requests of specified request type by the Branch Prediction unit. Specify one or more mask bits to qualify the request type(s).   |
|                   | ESCR restrictions        | MSR_BPU_ESCR0<br>MSR_BPU_ESCR1 |   |
|                   | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3     |   |
|                   | ESCR Event Select        | 03H                            | ESCR[31:25]   |
|                   | ESCR Event Mask          | Bit 0: TCMISS                  | ESCR[24:9]<br>Trace cache lookup miss   |
|                   | CCCR Select              | 00H                            | CCCR[15:13]   |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name      | Event Parameters         | Parameter Value                       | Description  |
|-----------------|--------------------------|---------------------------------------|--|
| ITLB_reference  |                          |                                       | This event counts translations using the Instruction Translation Look-aside Buffer (ITLB).   |
|                 | ESCR restrictions        | MSR_ITLB_ESCR0<br>MSR_ITLB_ESCR1      |  |
|                 | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3            |  |
|                 | ESCR Event Select        | 18H                                   | ESCR[31:25]  |
|                 | ESCR Event Mask          | Bit<br>0: HIT<br>1: MISS<br>2: HIT_UC | ESCR[24:9]<br><br>ITLB hit<br>ITLB miss<br>Uncacheable ITLB hit  |
|                 | CCCR Select              | 03H                                   | CCCR[15:13]  |
|                 | Event Specific Notes     |                                       | All page references regardless of the page size are looked up as actual 4-KByte pages. Use the page_walk_type event with the ITMISS mask for a more conservative count.                  |
| memory_cancel   |                          |                                       | This event counts the canceling of various type of request in the Data cache Address Control unit (DAC). Specify one or more mask bits to select the type of requests that are canceled. |
|                 | ESCR restrictions        | MSR_DAC_ESCR0<br>MSR_DAC_ESCR1        |  |
|                 | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11          |  |
|                 | ESCR Event Select        | 02H                                   | ESCR[31:25]  |
|                 | ESCR Event Mask          | Bit<br>2: ST_RB_FULL<br>3: 64K_CONF   | ESCR[24:9]<br><br>Replayed because no store request buffer is available.<br>Conflicts due to 64-KByte aliasing.  |
|                 | CCCR Select              | 05H                                   | CCCR[15:13]  |
|                 | Event Specific Notes     |                                       | All_CACHE_MISS includes uncacheable memory in count.   |
| memory_complete |                          |                                       | This event counts the completion of a load split, store split, uncacheable (UC) split, or UC load. Specify one or more mask bits to select the operations to be counted.                 |
|                 | ESCR restrictions        | MSR_SAAT_ESCR0<br>MSR_SAAT_ESCR1      |  |
|                 | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11          |  |
|                 | ESCR Event Select        | 08H                                   | ESCR[31:25]  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name        | Event Parameters         | Parameter Value                  | Description   |
|-------------------|--------------------------|----------------------------------|---|
|                   | ESCR Event Mask          | Bit<br>0: LSC<br>1: SSC          | ESCR[24:9]<br><br>Load split completed, excluding UC/WC loads.<br>Any split stores completed.   |
|                   | CCCR Select              | 02H                              | CCCR[15:13]   |
| load_port_replay  |                          |                                  | This event counts replayed events at the load port. Specify one or more mask bits to select the cause of the replay.  |
|                   | ESCR restrictions        | MSR_SAAT_ESCR0<br>MSR_SAAT_ESCR1 |   |
|                   | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |   |
|                   | ESCR Event Select        | 04H                              | ESCR[31:25]   |
|                   | ESCR Event Mask          | Bit 1: SPLIT_LD                  | ESCR[24:9]<br>Split load.   |
|                   | CCCR Select              | 02H                              | CCCR[15:13]   |
|                   | Event Specific Notes     |                                  | Must use ESCR1 for at-retirement counting.  |
| store_port_replay |                          |                                  | This event counts replayed events at the store port. Specify one or more mask bits to select the cause of the replay.   |
|                   | ESCR restrictions        | MSR_SAAT_ESCR0<br>MSR_SAAT_ESCR1 |   |
|                   | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |   |
|                   | ESCR Event Select        | 05H                              | ESCR[31:25]   |
|                   | ESCR Event Mask          | Bit 1: SPLIT_ST                  | ESCR[24:9]<br>Split store   |
|                   | CCCR Select              | 02H                              | CCCR[15:13]   |
|                   | Event Specific Notes     |                                  | Must use ESCR1 for at-retirement counting.  |
| MOB_load_replay   |                          |                                  | This event triggers if the memory order buffer (MOB) caused a load operation to be replayed. Specify one or more mask bits to select the cause of the replay. |
|                   | ESCR restrictions        | MSR_MOB_ESCR0<br>MSR_MOB_ESCR1   |   |
|                   | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3       |   |
|                   | ESCR Event Select        | 03H                              | ESCR[31:25]   |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name          | Event Parameters         | Parameter Value                       | Description   |
|---------------------|--------------------------|---------------------------------------|---|
|                     | ESCR Event Mask          | Bit<br>1: NO_STA<br><br>3: NO_STD     | ESCR[24:9]<br><br>Replayed because of unknown store address.<br>Replayed because of unknown store data.   |
|                     |                          | 4: PARTIAL_DATA<br><br>5: UNALGN_ADDR | Replayed because of partially overlapped data access between the load and store operations.<br>Replayed because the lower 4 bits of the linear address do not match between the load and store operations.  |
|                     | CCCR Select              | 02H                                   | CCCR[15:13]   |
| page_walk_type      |                          |                                       | This event counts various types of page walks that the page miss handler (PMH) performs.  |
|                     | ESCR restrictions        | MSR_PMH_ESCR0<br>MSR_PMH_ESCR1        |   |
|                     | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3            |   |
|                     | ESCR Event Select        | 01H                                   | ESCR[31:25]   |
|                     | ESCR Event Mask          | Bit<br>0: DTMISS<br><br>1: ITMISS     | ESCR[24:9]<br><br>Page walk for a data TLB miss (either load or store).<br>Page walk for an instruction TLB miss.   |
|                     | CCCR Select              | 04H                                   | CCCR[15:13]   |
| BSQ_cache_reference |                          |                                       | This event counts cache references (2nd level cache or 3rd level cache) as seen by the bus unit.<br><br>Specify one or more mask bit to select an access according to the access type (read type includes both load and RFO, write type includes writebacks and evictions) and the access result (hit, misses). |
|                     | ESCR restrictions        | MSR_BSU_ESCR0<br>MSR_BSU_ESCR1        |   |
|                     | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3            |   |
|                     | ESCR Event Select        | 0CH                                   | ESCR[31:25]   |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name     | Event Parameters     | Parameter Value   | Description   |
|----------------|----------------------|---|---|
|                |                      | Bit<br>0: RD_2ndL_HITS<br>1: RD_2ndL_HITE<br>2: RD_2ndL_HITM<br>3: RD_3rdL_HITS<br>4: RD_3rdL_HITE<br>5: RD_3rdL_HITM | ESCR[24:9]<br><br>Read 2nd level cache hit Shared (includes load and RFO).<br>Read 2nd level cache hit Exclusive (includes load and RFO).<br>Read 2nd level cache hit Modified (includes load and RFO).<br>Read 3rd level cache hit Shared (includes load and RFO).<br>Read 3rd level cache hit Exclusive (includes load and RFO).<br>Read 3rd level cache hit Modified (includes load and RFO).  |
|                | ESCR Event Mask      | 8: RD_2ndL_MISS<br>9: RD_3rdL_MISS<br>10: WR_2ndL_MISS  | Read 2nd level cache miss (includes load and RFO).<br>Read 3rd level cache miss (includes load and RFO).<br>A Writeback lookup from DAC misses the 2nd level cache (unlikely to happen).  |
|                | CCCR Select          | 07H   | CCCR[15:13]   |
|                | Event Specific Notes |   | 1: The implementation of this event in current Pentium 4 and Xeon processors treats either a load operation or a request for ownership (RFO) request as a “read” type operation.<br>2: Currently this event causes both over and undercounting by as much as a factor of two due to an erratum.<br>3: It is possible for a transaction that is started as a prefetch to change the transaction’s internal status, making it no longer a prefetch. or change the access result status (hit, miss) as seen by this event.   |
| IOQ_allocation |                      |   | This event counts the various types of transactions on the bus. A count is generated each time a transaction is allocated into the IOQ that matches the specified mask bits. An allocated entry can be a sector (64 bytes) or a chunks of 8 bytes.<br><br>Requests are counted once per retry. The event mask bits constitute 4 bit fields. A transaction type is specified by interpreting the values of each bit field.<br><br>Specify one or more event mask bits in a bit field to select the value of the bit field.<br><br>Each field (bits 0-4 are one field) are independent of and can be ORed with the others. The request type field is further combined with bit 5 and 6 to form a binary expression. Bits 7 and 8 form a bit field to specify the memory type of the target address.<br><br>Bits 13 and 14 form a bit field to specify the source agent of the request. Bit 15 affects read operation only. The event is triggered by evaluating the logical expression: (((Request type) OR Bit 5 OR Bit 6) OR (Memory type)) AND (Source agent). |



**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name | Event Parameters         | Parameter Value  | Description   |
|------------|--------------------------|--|---|
|            | ESCR restrictions        | MSR_FSB_ESCR0,<br>MSR_FSB_ESCR1  |   |
|            | Counter numbers per ESCR | ESCR0: 0, 1;<br>ESCR1: 2, 3  |   |
|            | ESCR Event Select        | 03H  | ESCR[31:25]   |
|            | ESCR Event Mask          | Bits<br>0-4 (single field)<br>5: ALL_READ<br>6: ALL_WRITE<br>7: MEM_UC<br>8: MEM_WC<br>9: MEM_WT<br>10: MEM_WP<br>11: MEM_WB<br>13: OWN<br><br>14: OTHER<br>15: PREFETCH | ESCR[24:9]<br><br>Bus request type (use 00001 for invalid or default).<br>Count read entries.<br>Count write entries.<br>Count UC memory access entries.<br>Count WC memory access entries.<br>Count write-through (WT) memory access entries.<br>Count write-protected (WP) memory access entries.<br>Count WB memory access entries.<br>Count all store requests driven by processor, as opposed to other processor or DMA.<br>Count all requests driven by other processors or DMA.<br>Include HW and SW prefetch requests in the count.   |
|            | CCCR Select              | 06H  | CCCR[15:13]   |
|            | Event Specific Notes     |  | <p>1: If PREFETCH bit is cleared, sectors fetched using prefetch are excluded in the counts. If PREFETCH bit is set, all sectors or chunks read are counted.</p> <p>2: Specify the edge trigger in CCCR to avoid double counting.</p> <p>3: The mapping of interpreted bit field values to transaction types may differ with different processor model implementations of the Pentium 4 processor family. Applications that program performance monitoring events should use CPUID to determine processor models when using this event. The logic equations that trigger the event are model-specific (see 4a and 4b below).</p> <p>4a: For Pentium 4 and Xeon Processors starting with CPUID Model field encoding equal to 2 or greater, this event is triggered by evaluating the logical expression ((Request type) and (Bit 5 or Bit 6) and (Memory type) and (Source agent)).</p> <p>4b: For Pentium 4 and Xeon Processors with CPUID Model field encoding less than 2, this event is triggered by evaluating the logical expression [((Request type) or Bit 5 or Bit 6) or (Memory type)] and (Source agent). Note that event mask bits for memory type are ignored if either ALL_READ or ALL_WRITE is specified.</p> <p>5: This event is known to ignore CPL in early implementations of Pentium 4 and Xeon Processors. Both user requests and OS requests are included in the count. This behavior is fixed starting with Pentium 4 and Xeon Processors with CPUID signature F27H (Family 15, Model 2, Stepping 7).</p> |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name         | Event Parameters         | Parameter Value   | Description   |
|--------------------|--------------------------|---|---|
|                    |                          |   | <p>6: For write-through (WT) and write-protected (WP) memory types, this event counts reads as the number of 64-byte sectors. Writes are counted by individual chunks.</p> <p>7: For uncacheable (UC) memory types, this event counts the number of 8-byte chunks allocated.</p> <p>8: For Pentium 4 and Xeon Processors with CPUID Signature less than F27H, only MSR_FSB_ESCR0 is available.</p>  |
| IOQ_active_entries |                          |   | <p>This event counts the number of entries (clipped at 15) in the IOQ that are active. An allocated entry can be a sector (64 bytes) or a chunks of 8 bytes.</p> <p>The event must be programmed in conjunction with IOQ_allocation. Specify one or more event mask bits to select the transactions that is counted.</p>  |
|                    | ESCR restrictions        | MSR_FSB_ESCR1   |   |
|                    | Counter numbers per ESCR | ESCR1: 2, 3   |   |
|                    | ESCR Event Select        | 01AH  | ESCR[30:25]   |
|                    | ESCR Event Mask          | Bits<br>0-4 (single field)<br>5: ALL_READ<br>6: ALL_WRITE<br>7: MEM_UC<br>8: MEM_WC<br>9: MEM_WT<br>10: MEM_WP<br>11: MEM_WB<br>13: O/WN<br>14: OTHER<br>15: PREFETCH | ESCR[24:9]<br><br>Bus request type (use 00001 for invalid or default).<br>Count read entries.<br>Count write entries.<br>Count UC memory access entries.<br>Count WC memory access entries.<br>Count write-through (WT) memory access entries.<br>Count write-protected (WP) memory access entries.<br>Count WB memory access entries.<br>Count all store requests driven by processor, as opposed to other processor or DMA.<br>Count all requests driven by other processors or DMA.<br>Include HW and SW prefetch requests in the count.   |
|                    | CCCR Select              | 06H   | CCCR[15:13]   |
|                    | Event Specific Notes     |   | <p>1: Specified desired mask bits in ESCR0 and ESCR1.</p> <p>2: See the ioq_allocation event for descriptions of the mask bits.</p> <p>3: Edge triggering should not be used when counting cycles.</p> <p>4: The mapping of interpreted bit field values to transaction types may differ across different processor model implementations of the Pentium 4 processor family. Applications that programs performance monitoring events should use the CPUID instruction to detect processor models when using this event. The logical expression that triggers this event as describe below:</p> <p>5a: For Pentium 4 and Xeon Processors starting with CPUID MODEL field encoding equal to 2 or greater, this event is triggered by evaluating the logical expression ((Request type) and (Bit 5 or Bit 6) and (Memory type) and (Source agent)).</p> |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name        | Event Parameters         | Parameter Value  | Description   |
|-------------------|--------------------------|--|---|
|                   |                          |  | <p>5b: For Pentium 4 and Xeon Processors starting with CPUID MODEL field encoding less than 2, this event is triggered by evaluating the logical expression [((Request type) or Bit 5 or Bit 6) or (Memory type)] and (Source agent). Event mask bits for memory type are ignored if either ALL_READ or ALL_WRITE is specified.</p> <p>5c: This event is known to ignore CPL in the current implementations of Pentium 4 and Xeon Processors Both user requests and OS requests are included in the count.</p> <p>6: An allocated entry can be a full line (64 bytes) or in individual chunks of 8 bytes.</p>   |
| FSB_data_activity |                          |  | This event increments once for each DRDY or DBSY event that occurs on the front side bus. The event allows selection of a specific DRDY or DBSY event.  |
|                   | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1   |   |
|                   | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3   |   |
|                   | ESCR Event Select        | 17H  | ESCR[31:25]   |
|                   | ESCR Event Mask          | <p>Bit 0:<br/>DRDY_DRV</p> <p>1: DRDY_OWN</p> <p>2: DRDY_OTHER</p> <p>3: DBSY_DRV</p> <p>4: DBSY_OWN</p> | <p>ESCR[24:9]</p> <p>Count when this processor drives data onto the bus - includes writes and implicit writebacks.<br/>Asserted two processor clock cycles for partial writes and 4 processor clocks (usually in consecutive bus clocks) for full line writes.</p> <p>Count when this processor reads data from the bus - includes loads and some PIC transactions. Asserted two processor clock cycles for partial reads and 4 processor clocks (usually in consecutive bus clocks) for full line reads.<br/>Count DRDY events that we drive.<br/>Count DRDY events sampled that we own.</p> <p>Count when data is on the bus but not being sampled by the processor. It may or may not be being driven by this processor.<br/>Asserted two processor clock cycles for partial transactions and 4 processor clocks (usually in consecutive bus clocks) for full line transactions.</p> <p>Count when this processor reserves the bus for use in the next bus cycle in order to drive data. Asserted for two processor clock cycles for full line writes and not at all for partial line writes.<br/>May be asserted multiple times (in consecutive bus clocks) if we stall the bus waiting for a cache lock to complete.</p> <p>Count when some agent reserves the bus for use in the next bus cycle to drive data that this processor will sample.<br/>Asserted for two processor clock cycles for full line writes and not at all for partial line writes. May be asserted multiple times (all one bus clock apart) if we stall the bus for some reason.</p> |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name     | Event Parameters         | Parameter Value  | Description  |
|----------------|--------------------------|--|--|
|                |                          | 5:DBSY_OTHER   | Count when some agent reserves the bus for use in the next bus cycle to drive data that this processor will NOT sample. It may or may not be being driven by this processor.<br><br>Asserted two processor clock cycles for partial transactions and 4 processor clocks (usually in consecutive bus clocks) for full line transactions.  |
|                | CCCR Select              | 06H  | CCCR[15:13]  |
|                | Event Specific Notes     |  | Specify edge trigger in the CCCR MSR to avoid double counting.<br>DRDY_OWN and DRDY_OTHER are mutually exclusive; similarly for DBSY_OWN and DBSY_OTHER.   |
| BSQ_allocation |                          |  | This event counts allocations in the Bus Sequence Unit (BSQ) according to the specified mask bit encoding. The event mask bits consist of four sub-groups: <ul style="list-style-type: none"> <li>▪ Request type.</li> <li>▪ Request length.</li> <li>▪ Memory type.</li> <li>▪ Sub-group consisting mostly of independent bits (bits 5, 6, 7, 8, 9, and 10).</li> </ul> Specify an encoding for each sub-group.   |
|                | ESCR restrictions        | MSR_BSU_ESCR0  |  |
|                | Counter numbers per ESCR | ESCR0: 0, 1  |  |
|                | ESCR Event Select        | 05H  | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit<br>0: REQ_TYPE0<br>1: REQ_TYPE1<br><br>2: REQ_LEN0<br>3: REQ_LEN1<br><br>5: REQ_IO_TYPE<br>6: REQ_LOCK_TYPE<br>7: REQ_CACHE_TYPE<br>8: REQ_SPLIT_TYPE<br>9: REQ_DEM_TYPE<br><br>10: REQ_ORD_TYPE | ESCR[24:9]<br><br>Request type encoding (bit 0 and 1) are:<br>0 - Read (excludes read invalidate).<br>1 - Read invalidate.<br>2 - Write (other than writebacks).<br>3 - Writeback (evicted from cache). (public)<br><br>Request length encoding (bit 2, 3) are:<br>0 - 0 chunks<br>1 - 1 chunks<br>3 - 8 chunks<br><br>Request type is input or output.<br>Request type is bus lock.<br><br>Request type is cacheable.<br><br>Request type is a bus 8-byte chunk split across 8-byte boundary.<br><br>Request type is a demand if set. Request type is HW.SW prefetch if 0.<br><br>Request is an ordered type. |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name         | Event Parameters         | Parameter Value                                 | Description  |
|--------------------|--------------------------|---|--|
|                    |                          | 11: MEM_TYPE0<br>12: MEM_TYPE1<br>13: MEM_TYPE2 | Memory type encodings (bit 11-13) are:<br>0 - UC<br>1 - WC<br>4 - WT<br>5 - WP<br>6 - WB   |
|                    | CCCR Select              | 07H   | CCCR[15:13]  |
|                    | Event Specific Notes     |   | 1: Specify edge trigger in CCCR to avoid double counting.<br>2: A writebacks to 3rd level cache from 2nd level cache counts as a separate entry, this is in addition to the entry allocated for a request to the bus.<br>3: A read request to WB memory type results in a request to the 64-byte sector, containing the target address, followed by a prefetch request to an adjacent sector.<br>4: For Pentium 4 and Xeon processors with CPUID model encoding value equals to 0 and 1, an allocated BSQ entry includes both the demand sector and prefetched 2nd sector.<br>5: An allocated BSQ entry for a data chunk is any request less than 64 bytes.<br>6a: This event may undercount for requests of split type transactions if the data address straddled across modulo-64 byte boundary.<br>6b: This event may undercount for requests of read request of 16-byte operands from WC or UC address.<br>6c: This event may undercount WC partial requests originated from store operands that are dwords. |
| bsq_active_entries |                          |   | This event represents the number of BSQ entries (clipped at 15) currently active (valid) which meet the subevent mask criteria during allocation in the BSQ. Active request entries are allocated on the BSQ until de-allocated.<br>De-allocation of an entry does not necessarily imply the request is filled. This event must be programmed in conjunction with BSQ_allocation. Specify one or more event mask bits to select the transactions that is counted.  |
|                    | ESCR restrictions        | ESCR1   |  |
|                    | Counter numbers per ESCR | ESCR1: 2, 3                                     |  |
|                    | ESCR Event Select        | 06H   | ESCR[30:25]  |
|                    | ESCR Event Mask          |   | ESCR[24:9]   |
|                    | CCCR Select              | 07H   | CCCR[15:13]  |
|                    | Event Specific Notes     |   | 1: Specified desired mask bits in ESCR0 and ESCR1.<br>2: See the BSQ_allocation event for descriptions of the mask bits.<br>3: Edge triggering should not be used when counting cycles.<br>4: This event can be used to estimate the latency of a transaction from allocation to de-allocation in the BSQ. The latency observed by BSQ_allocation includes the latency of FSB, plus additional overhead.   |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name       | Event Parameters         | Parameter Value                  | Description  |
|------------------|--------------------------|----------------------------------|--|
|                  |                          |                                  | <p>5: Additional overhead may include the time it takes to issue two requests (the sector by demand and the adjacent sector via prefetch). Since adjacent sector prefetches have lower priority than demand fetches, on a heavily used system there is a high probability that the adjacent sector prefetch will have to wait until the next bus arbitration.</p> <p>6: For Pentium 4 and Xeon processors with CPUID model encoding value less than 3, this event is updated every clock.</p> <p>7: For Pentium 4 and Xeon processors with CPUID model encoding value equals to 3 or 4, this event is updated every other clock.</p>   |
| SSE_input_assist |                          |                                  | This event counts the number of times an assist is requested to handle problems with input operands for SSE/SSE2/SSE3 operations; most notably denormal source operands when the DAZ bit is not set. Set bit 15 of the event mask to use this event.   |
|                  | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|                  | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |  |
|                  | ESCR Event Select        | 34H                              | ESCR[31:25]  |
|                  | ESCR Event Mask          | 15: ALL                          | ESCR[24:9]<br>Count assists for SSE/SSE2/SSE3 $\mu$ ops.   |
|                  | CCCR Select              | 01H                              | CCCR[15:13]  |
|                  | Event Specific Notes     |                                  | <p>1: Not all requests for assists are actually taken. This event is known to overcount in that it counts requests for assists from instructions on the non-retired path that do not incur a performance penalty. An assist is actually taken only for non-bogus <math>\mu</math>ops. Any appreciable counts for this event are an indication that the DAZ or FTZ bit should be set and/or the source code should be changed to eliminate the condition.</p> <p>2: Two common situations for an SSE/SSE2/SSE3 operation needing an assist are: (1) when a denormal constant is used as an input and the Denormals-Are-Zero (DAZ) mode is not set, (2) when the input operand uses the underflowed result of a previous SSE/SSE2/SSE3 operation and neither the DAZ nor Flush-To-Zero (FTZ) modes are set.</p> <p>3: Enabling the DAZ mode prevents SSE/SSE2/SSE3 operations from needing assists in the first situation. Enabling the FTZ mode prevents SSE/SSE2/SSE3 operations from needing assists in the second situation.</p> |
| packed_SP_uop    |                          |                                  | This event increments for each packed single-precision $\mu$ op, specified through the event mask for detection.   |
|                  | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|                  | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |  |
|                  | ESCR Event Select        | 08H                              | ESCR[31:25]  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name    | Event Parameters         | Parameter Value                  | Description  |
|---------------|--------------------------|----------------------------------|--|
|               | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on packed single-precision operands.   |
|               | CCCR Select              | 01H                              | CCCR[15:13]  |
|               | Event Specific Notes     |                                  | 1: If an instruction contains more than one packed SP $\mu$ ops, each packed SP $\mu$ op that is specified by the event mask will be counted.<br>2: This metric counts instances of packed memory $\mu$ ops in a repeat move string. |
| packed_DP_uop |                          |                                  | This event increments for each packed double-precision $\mu$ op, specified through the event mask for detection.   |
|               | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|               | Counter numbers per ESCR | ESCRO: 8, 9<br>ESCR1: 10, 11     |  |
|               | ESCR Event Select        | OCH                              | ESCR[31:25]  |
|               | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on packed double-precision operands.   |
|               | CCCR Select              | 01H                              | CCCR[15:13]  |
|               | Event Specific Notes     |                                  | If an instruction contains more than one packed DP $\mu$ ops, each packed DP $\mu$ op that is specified by the event mask will be counted.   |
| scalar_SP_uop |                          |                                  | This event increments for each scalar single-precision $\mu$ op, specified through the event mask for detection.   |
|               | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|               | Counter numbers per ESCR | ESCRO: 8, 9<br>ESCR1: 10, 11     |  |
|               | ESCR Event Select        | OAH                              | ESCR[31:25]  |
|               | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on scalar single-precision operands.   |
|               | CCCR Select              | 01H                              | CCCR[15:13]  |
|               | Event Specific Notes     |                                  | If an instruction contains more than one scalar SP $\mu$ ops, each scalar SP $\mu$ op that is specified by the event mask will be counted.   |
| scalar_DP_uop |                          |                                  | This event increments for each scalar double-precision $\mu$ op, specified through the event mask for detection.   |
|               | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|               | Counter numbers per ESCR | ESCRO: 8, 9<br>ESCR1: 10, 11     |  |
|               | ESCR Event Select        | 0EH                              | ESCR[31:25]  |
|               | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on scalar double-precision operands.   |
|               | CCCR Select              | 01H                              | CCCR[15:13]  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name     | Event Parameters         | Parameter Value                  | Description  |
|----------------|--------------------------|----------------------------------|--|
|                | Event Specific Notes     |                                  | If an instruction contains more than one scalar DP $\mu$ ops, each scalar DP $\mu$ op that is specified by the event mask is counted.          |
| 64bit_MMX_uop  |                          |                                  | This event increments for each MMX instruction, which operate on 64-bit SIMD operands.   |
|                | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|                | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |  |
|                | ESCR Event Select        | 02H                              | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on 64-bit SIMD integer operands in memory or MMX registers.  |
|                | CCCR Select              | 01H                              | CCCR[15:13]  |
|                | Event Specific Notes     |                                  | If an instruction contains more than one 64-bit MMX $\mu$ ops, each 64-bit MMX $\mu$ op that is specified by the event mask will be counted.   |
| 128bit_MMX_uop |                          |                                  | This event increments for each integer SIMD SSE2 instruction, which operate on 128-bit SIMD operands.  |
|                | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|                | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |  |
|                | ESCR Event Select        | 1AH                              | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all $\mu$ ops operating on 128-bit SIMD integer operands in memory or XMM registers.                                       |
|                | CCCR Select              | 01H                              | CCCR[15:13]  |
|                | Event Specific Notes     |                                  | If an instruction contains more than one 128-bit MMX $\mu$ ops, each 128-bit MMX $\mu$ op that is specified by the event mask will be counted. |
| x87_FP_uop     |                          |                                  | This event increments for each x87 floating-point $\mu$ op, specified through the event mask for detection.                                    |
|                | ESCR restrictions        | MSR_FIRM_ESCRO<br>MSR_FIRM_ESCR1 |  |
|                | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11     |  |
|                | ESCR Event Select        | 04H                              | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit 15: ALL                      | ESCR[24:9]<br>Count all x87 FP $\mu$ ops.  |
|                | CCCR Select              | 01H                              | CCCR[15:13]  |



**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name          | Event Parameters         | Parameter Value                | Description  |
|---------------------|--------------------------|--------------------------------|--|
|                     | Event Specific Notes     |                                | 1: If an instruction contains more than one x87 FP $\mu$ ops, each x87 FP $\mu$ op that is specified by the event mask will be counted.<br>2: This event does not count x87 FP $\mu$ op for load, store, move between registers. |
| TC_misc             |                          |                                | This event counts miscellaneous events detected by the TC. The counter will count twice for each occurrence.   |
|                     | ESCR restrictions        | MSR_TC_ESCR0<br>MSR_TC_ESCR1   |  |
|                     | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7     |  |
|                     | ESCR Event Select        | 06H                            | ESCR[31:25]  |
|                     | CCCR Select              | 01H                            | CCCR[15:13]  |
|                     | ESCR Event Mask          | Bit 4: FLUSH                   | ESCR[24:9]<br>Number of flushes  |
| global_power_events |                          |                                | This event accumulates the time during which a processor is not stopped.   |
|                     | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1 |  |
|                     | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3     |  |
|                     | ESCR Event Select        | 013H                           | ESCR[31:25]  |
|                     | ESCR Event Mask          | Bit 0: Running                 | ESCR[24:9]<br>The processor is active (includes the handling of HLT STPCLK and throttling.   |
|                     | CCCR Select              | 06H                            | CCCR[15:13]  |
| tc_ms_xfer          |                          |                                | This event counts the number of times that uop delivery changed from TC to MS ROM.   |
|                     | ESCR restrictions        | MSR_MS_ESCR0<br>MSR_MS_ESCR1   |  |
|                     | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7     |  |
|                     | ESCR Event Select        | 05H                            | ESCR[31:25]  |
|                     | ESCR Event Mask          | Bit 0: CISC                    | ESCR[24:9]<br>A TC to MS transfer occurred.  |
|                     | CCCR Select              | 0H                             | CCCR[15:13]  |
| uop_queue_writes    |                          |                                | This event counts the number of valid uops written to the uop queue. Specify one or more mask bits to select the source type of writes.  |
|                     | ESCR restrictions        | MSR_MS_ESCR0<br>MSR_MS_ESCR1   |  |
|                     | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7     |  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name                  | Event Parameters         | Parameter Value  | Description  |
|-----------------------------|--------------------------|--|--|
|                             | ESCR Event Select        | 09H  | ESCR[31:25]  |
|                             | ESCR Event Mask          | Bit<br>0: FROM_TC_BUILD<br>1: FROM_TC_DELIVER<br>2: FROM_ROM | ESCR[24:9]<br><br>The uops being written are from TC build mode.<br><br>The uops being written are from TC deliver mode.<br>The uops being written are from microcode ROM.   |
|                             | CCCR Select              | 0H   | CCCR[15:13]  |
| retired_mispred_branch_type |                          |  | This event counts retiring mispredicted branches by type.  |
|                             | ESCR restrictions        | MSR_TBPU_ESCR0<br>MSR_TBPU_ESCR1                             |  |
|                             | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7                                   |  |
|                             | ESCR Event Select        | 05H  | ESCR[30:25]  |
|                             | ESCR Event Mask          | Bit<br>1: CONDITIONAL<br>2: CALL                             | ESCR[24:9]<br><br>Conditional jumps.<br>Indirect call branches.  |
|                             |                          | 3: RETURN<br>4: INDIRECT                                     | Return branches.<br>Returns, indirect calls, or indirect jumps.  |
|                             | CCCR Select              | 02H  | CCCR[15:13]  |
|                             | Event Specific Notes     |  | This event may overcount conditional branches if: <ul style="list-style-type: none"> <li>▪ Mispredictions cause the trace cache and delivery engine to build new traces.</li> <li>▪ When the processor's pipeline is being cleared.</li> </ul> |
| retired_branch_type         |                          |  | This event counts retiring branches by type. Specify one or more mask bits to qualify the branch by its type.  |
|                             | ESCR restrictions        | MSR_TBPU_ESCR0<br>MSR_TBPU_ESCR1                             |  |
|                             | Counter numbers per ESCR | ESCR0: 4, 5<br>ESCR1: 6, 7                                   |  |
|                             | ESCR Event Select        | 04H  | ESCR[30:25]  |
|                             | ESCR Event Mask          | Bit<br>1: CONDITIONAL<br>2: CALL<br>3: RETURN<br>4: INDIRECT | ESCR[24:9]<br><br>Conditional jumps.<br>Direct or indirect calls.<br>Return branches.<br>Returns, indirect calls, or indirect jumps.   |
|                             | CCCR Select              | 02H  | CCCR[15:13]  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name           | Event Parameters         | Parameter Value  | Description   |
|----------------------|--------------------------|--|---|
|                      | Event Specific Notes     |  | This event may overcount conditional branches if : <ul style="list-style-type: none"> <li>▪ Mispredictions cause the trace cache and delivery engine to build new traces.</li> <li>▪ When the processor’s pipeline is being cleared.</li> </ul> |
| resource_stall       |                          |  | This event monitors the occurrence or latency of stalls in the Allocator.   |
|                      | ESCR restrictions        | MSR_ALF_ESCR0<br>MSR_ALF_ESCR1   |   |
|                      | Counter numbers per ESCR | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17   |   |
|                      | ESCR Event Select        | 01H  | ESCR[30:25]   |
|                      | Event Masks              | Bit<br>5: SBFULL   | ESCR[24:9]<br><br>A Stall due to lack of store buffers.   |
|                      | CCCR Select              | 01H  | CCCR[15:13]   |
|                      | Event Specific Notes     |  | This event may not be supported in all models of the processor family.  |
| WC_Buffer            |                          |  | This event counts Write Combining Buffer operations that are selected by the event mask.  |
|                      | ESCR restrictions        | MSR_DAC_ESCR0<br>MSR_DAC_ESCR1   |   |
|                      | Counter numbers per ESCR | ESCR0: 8, 9<br>ESCR1: 10, 11   |   |
|                      | ESCR Event Select        | 05H  | ESCR[30:25]   |
|                      | Event Masks              | Bit<br>0: WCB_EVICTS   | ESCR[24:9]<br><br>WC Buffer evictions of all causes.  |
|                      |                          | 1: WCB_FULL_EVICT  | WC Buffer eviction: no WC buffer is available.  |
|                      | CCCR Select              | 05H  | CCCR[15:13]   |
| Event Specific Notes |                          | This event is useful for detecting the subset of 64K aliasing cases that are more costly (i.e. 64K aliasing cases involving stores) as long as there are no significant contributions due to write combining buffer full or hit-modified conditions. |   |
| b2b_cycles           |                          |  | This event can be configured to count the number back-to-back bus cycles using sub-event mask bits 1 through 6.   |
|                      | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1   |   |
|                      | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3   |   |
|                      | ESCR Event Select        | 016H   | ESCR[30:25]   |
|                      | Event Masks              | Bit  | ESCR[24:9]  |

**Table 19-31. Performance Monitoring Events Supported by Intel NetBurst® Microarchitecture for Non-Retirement Counting (Contd.)**

| Event Name | Event Parameters         | Parameter Value                | Description   |
|------------|--------------------------|--------------------------------|---|
|            | CCCR Select              | 03H                            | CCCR[15:13]   |
|            | Event Specific Notes     |                                | This event may not be supported in all models of the processor family.                                      |
| bnr        |                          |                                | This event can be configured to count bus not ready conditions using sub-event mask bits 0 through 2.       |
|            | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1 |   |
|            | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3     |   |
|            | ESCR Event Select        | 08H                            | ESCR[30:25]   |
|            | Event Masks              | Bit                            | ESCR[24:9]  |
|            | CCCR Select              | 03H                            | CCCR[15:13]   |
|            | Event Specific Notes     |                                | This event may not be supported in all models of the processor family.                                      |
| snoop      |                          |                                | This event can be configured to count snoop hit modified bus traffic using sub-event mask bits 2, 6 and 7.  |
|            | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1 |   |
|            | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3     |   |
|            | ESCR Event Select        | 06H                            | ESCR[30:25]   |
|            | Event Masks              | Bit                            | ESCR[24:9]  |
|            | CCCR Select              | 03H                            | CCCR[15:13]   |
|            | Event Specific Notes     |                                | This event may not be supported in all models of the processor family.                                      |
| Response   |                          |                                | This event can be configured to count different types of responses using sub-event mask bits 1,2, 8, and 9. |
|            | ESCR restrictions        | MSR_FSB_ESCR0<br>MSR_FSB_ESCR1 |   |
|            | Counter numbers per ESCR | ESCR0: 0, 1<br>ESCR1: 2, 3     |   |
|            | ESCR Event Select        | 04H                            | ESCR[30:25]   |
|            | Event Masks              | Bit                            | ESCR[24:9]  |
|            | CCCR Select              | 03H                            | CCCR[15:13]   |
|            | Event Specific Notes     |                                | This event may not be supported in all models of the processor family.                                      |

**Table 19-32. Performance Monitoring Events For Intel NetBurst® Microarchitecture for At-Retirement Counting**

| Event Name      | Event Parameters                    | Parameter Value   | Description  |
|-----------------|-------------------------------------|---|--|
| front_end_event |                                     |   | This event counts the retirement of tagged $\mu$ ops, which are specified through the front-end tagging mechanism. The event mask specifies bogus or non-bogus $\mu$ ops.  |
|                 | ESCR restrictions                   | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3  |  |
|                 | Counter numbers per ESCR            | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17  |  |
|                 | ESCR Event Select                   | 08H   | ESCR[31:25]  |
|                 | ESCR Event Mask                     | Bit<br>0: NBOGUS<br>1: BOGUS  | ESCR[24:9]<br><br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are bogus.   |
|                 | CCCR Select                         | 05H   | CCCR[15:13]  |
|                 | Can Support PEBS                    | Yes   |  |
|                 | Require Additional MSRs for tagging | Selected ESCRs and/or MSR_TC_PRECISE_EVENT  | See list of metrics supported by Front_end tagging in Table A-3  |
| execution_event |                                     |   | This event counts the retirement of tagged $\mu$ ops, which are specified through the execution tagging mechanism.<br><br>The event mask allows from one to four types of $\mu$ ops to be specified as either bogus or non-bogus $\mu$ ops to be tagged.   |
|                 | ESCR restrictions                   | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3  |  |
|                 | Counter numbers per ESCR            | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17  |  |
|                 | ESCR Event Select                   | 0CH   | ESCR[31:25]  |
|                 | ESCR Event Mask                     | Bit<br>0: NBOGUS0<br>1: NBOGUS1<br>2: NBOGUS2<br>3: NBOGUS3<br>4: BOGUS0<br>5: BOGUS1<br>6: BOGUS2<br>7: BOGUS3 | ESCR[24:9]<br><br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are bogus.<br>The marked $\mu$ ops are bogus.<br>The marked $\mu$ ops are bogus.<br>The marked $\mu$ ops are bogus. |
|                 | CCCR Select                         | 05H   | CCCR[15:13]  |
|                 | Event Specific Notes                |   | Each of the 4 slots to specify the bogus/non-bogus $\mu$ ops must be coordinated with the 4 TagValue bits in the ESCR (for example, NBOGUS0 must accompany a '1' in the lowest bit of the TagValue field in ESCR, NBOGUS1 must accompany a '1' in the next but lowest bit of the TagValue field).                      |
|                 | Can Support PEBS                    | Yes   |  |

**Table 19-32. Performance Monitoring Events For Intel NetBurst® Microarchitecture for At-Retirement Counting (Contd.)**

| Event Name    | Event Parameters                    | Parameter Value   | Description  |
|---------------|-------------------------------------|---|--|
|               | Require Additional MSRs for tagging | An ESCR for an upstream event                                       | See list of metrics supported by execution tagging in Table A-4.   |
| replay_event  |                                     |   | This event counts the retirement of tagged $\mu$ ops, which are specified through the replay tagging mechanism. The event mask specifies bogus or non-bogus $\mu$ ops.   |
|               | ESCR restrictions                   | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3                                      |  |
|               | Counter numbers per ESCR            | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17                              |  |
|               | ESCR Event Select                   | 09H   | ESCR[31:25]  |
|               | ESCR Event Mask                     | Bit<br>0: NBOGUS<br>1: BOGUS  | ESCR[24:9]<br><br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are bogus.   |
|               | CCCR Select                         | 05H   | CCCR[15:13]  |
|               | Event Specific Notes                |   | Supports counting tagged $\mu$ ops with additional MSRs.   |
|               | Can Support PEBS                    | Yes   |  |
|               | Require Additional MSRs for tagging | IA32_PEBS_ENABLE<br>MSR_PEBS_MATRIX_VERT<br>Selected ESCR           | See list of metrics supported by replay tagging in Table A-5.  |
| instr_retired |                                     |   | This event counts instructions that are retired during a clock cycle. Mask bits specify bogus or non-bogus (and whether they are tagged using the front-end tagging mechanism).  |
|               | ESCR restrictions                   | MSR_CRU_ESCR0<br>MSR_CRU_ESCR1                                      |  |
|               | Counter numbers per ESCR            | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17                              |  |
|               | ESCR Event Select                   | 02H   | ESCR[31:25]  |
|               | ESCR Event Mask                     | Bit<br>0: NBOGUSNTAG<br>1: NBOGUSTAG<br>2: BOGUSNTAG<br>3: BOGUSTAG | ESCR[24:9]<br><br>Non-bogus instructions that are not tagged.<br>Non-bogus instructions that are tagged.<br><br>Bogus instructions that are not tagged.<br>Bogus instructions that are tagged.   |
|               | CCCR Select                         | 04H   | CCCR[15:13]  |
|               | Event Specific Notes                |   | 1: The event count may vary depending on the microarchitectural states of the processor when the event detection is enabled.<br>2: The event may count more than once for some instructions with complex uop flows and were interrupted before retirement. |

**Table 19-32. Performance Monitoring Events For Intel NetBurst® Microarchitecture for At-Retirement Counting (Contd.)**

| Event Name     | Event Parameters         | Parameter Value                        | Description  |
|----------------|--------------------------|--|--|
|                | Can Support PEBS         | No                                     |  |
| uops_retired   |                          |  | This event counts $\mu$ ops that are retired during a clock cycle. Mask bits specify bogus or non-bogus.   |
|                | ESCR restrictions        | MSR_CRU_ESCR0<br>MSR_CRU_ESCR1         |  |
|                | Counter numbers per ESCR | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17 |  |
|                | ESCR Event Select        | 01H                                    | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit<br>0: NBOGUS<br>1: BOGUS           | ESCR[24:9]<br><br>The marked $\mu$ ops are not bogus.<br>The marked $\mu$ ops are bogus.   |
|                | CCCR Select              | 04H                                    | CCCR[15:13]  |
|                | Event Specific Notes     |  | P6: EMON_UOPS_RETIRE   |
|                | Can Support PEBS         | No                                     |  |
| uop_type       |                          |  | This event is used in conjunction with the front-end at-retirement mechanism to tag load and store $\mu$ ops.  |
|                | ESCR restrictions        | MSR_RAT_ESCR0<br>MSR_RAT_ESCR1         |  |
|                | Counter numbers per ESCR | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17 |  |
|                | ESCR Event Select        | 02H                                    | ESCR[31:25]  |
|                | ESCR Event Mask          | Bit<br>1: TAGLOADS<br>2: TAGSTORES     | ESCR[24:9]<br><br>The $\mu$ op is a load operation.<br>The $\mu$ op is a store operation.  |
|                | CCCR Select              | 02H                                    | CCCR[15:13]  |
|                | Event Specific Notes     |  | Setting the TAGLOADS and TAGSTORES mask bits does not cause a counter to increment. They are only used to tag uops.                                    |
|                | Can Support PEBS         | No                                     |  |
| branch_retired |                          |  | This event counts the retirement of a branch. Specify one or more mask bits to select any combination of taken, not-taken, predicted and mispredicted. |
|                | ESCR restrictions        | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3         | See Table 18-70 for the addresses of the ESCR MSRs   |
|                | Counter numbers per ESCR | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17 | The counter numbers associated with each ESCR are provided. The performance counters and corresponding CCCRs can be obtained from Table 18-70.         |
|                | ESCR Event Select        | 06H                                    | ESCR[31:25]  |

**Table 19-32. Performance Monitoring Events For Intel NetBurst® Microarchitecture for At-Retirement Counting (Contd.)**

| Event Name             | Event Parameters         | Parameter Value  | Description  |
|------------------------|--------------------------|--|--|
|                        | ESCR Event Mask          | Bit<br>0: MMNP<br>1: MMNM<br>2: MMTP<br>3: MMTM            | ESCR[24:9]<br><br>Branch not-taken predicted<br>Branch not-taken mispredicted<br>Branch taken predicted<br>Branch taken mispredicted                                 |
|                        | CCCR Select              | 05H  | CCCR[15:13]  |
|                        | Event Specific Notes     |  | P6: EMON_BR_INST_RETIRED   |
|                        | Can Support PEBS         | No   |  |
|                        |                          |  |  |
| mispred_branch_retired |                          |  | This event represents the retirement of mispredicted branch instructions.  |
|                        | ESCR restrictions        | MSR_CRU_ESCR0<br>MSR_CRU_ESCR1                             |  |
|                        | Counter numbers per ESCR | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17                     |  |
|                        | ESCR Event Select        | 03H  | ESCR[31:25]  |
|                        | ESCR Event Mask          | Bit 0: NBOGUS  | ESCR[24:9]<br>The retired instruction is not bogus.  |
|                        | CCCR Select              | 04H  | CCCR[15:13]  |
|                        | Can Support PEBS         | No   |  |
| x87_assist             |                          |  | This event counts the retirement of x87 instructions that required special handling.<br>Specifies one or more event mask bits to select the type of assistance.      |
|                        | ESCR restrictions        | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3                             |  |
|                        | Counter numbers per ESCR | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17                     |  |
|                        | ESCR Event Select        | 03H  | ESCR[31:25]  |
|                        | ESCR Event Mask          | Bit<br>0: FPSU<br>1: FPSO<br>2: POAO<br>3: POAU<br>4: PREA | ESCR[24:9]<br><br>Handle FP stack underflow.<br>Handle FP stack overflow.<br>Handle x87 output overflow.<br>Handle x87 output underflow.<br>Handle x87 input assist. |
|                        | CCCR Select              | 05H  | CCCR[15:13]  |
|                        | Can Support PEBS         | No   |  |



**Table 19-32. Performance Monitoring Events For Intel NetBurst® Microarchitecture for At-Retirement Counting (Contd.)**

| Event Name    | Event Parameters         | Parameter Value                                     | Description   |
|---------------|--------------------------|---|---|
| machine_clear |                          |   | This event increments according to the mask bit specified while the entire pipeline of the machine is cleared. Specify one of the mask bit to select the cause.   |
|               | ESCR restrictions        | MSR_CRU_ESCR2<br>MSR_CRU_ESCR3                      |   |
|               | Counter numbers per ESCR | ESCR2: 12, 13, 16<br>ESCR3: 14, 15, 17              |   |
|               | ESCR Event Select        | 02H   | ESCR[31:25]   |
|               | ESCR Event Mask          | Bit<br>0: CLEAR<br><br>2: MOCLEAR<br><br>6: SMCLEAR | ESCR[24:9]<br><br>Counts for a portion of the many cycles while the machine is cleared for any cause. Use Edge triggering for this bit only to get a count of occurrence versus a duration.<br><br>Increments each time the machine is cleared due to memory ordering issues.<br><br>Increments each time the machine is cleared due to self-modifying code issues. |
|               | CCCR Select              | 05H   | CCCR[15:13]   |
|               | Can Support PEBS         | No  |   |

**Table 19-33. Intel NetBurst® Microarchitecture Model-Specific Performance Monitoring Events (For Model Encoding 3, 4 or 6)**

| Event Name      | Event Parameters         | Parameter Value                        | Description  |
|-----------------|--------------------------|--|--|
| instr_completed |                          |  | This event counts instructions that have completed and retired during a clock cycle. Mask bits specify whether the instruction is bogus or non-bogus and whether they are: |
|                 | ESCR restrictions        | MSR_CRU_ESCR0<br>MSR_CRU_ESCR1         |  |
|                 | Counter numbers per ESCR | ESCR0: 12, 13, 16<br>ESCR1: 14, 15, 17 |  |
|                 | ESCR Event Select        | 07H                                    | ESCR[31:25]  |
|                 | ESCR Event Mask          | Bit<br>0: NBOGUS<br>1: BOGUS           | ESCR[24:9]<br><br>Non-bogus instructions<br>Bogus instructions   |
|                 | CCCR Select              | 04H                                    | CCCR[15:13]  |
|                 | Event Specific Notes     |  | This metric differs from instr_retired, since it counts instructions completed, rather than the number of times that instructions started.                                 |
|                 | Can Support PEBS         | No                                     |  |

**Table 19-34. List of Metrics Available for Front\_end Tagging (For Front\_end Event Only)**

| Front-end metric <sup>1</sup> | MSR_TC_PRECISE_EVENT MSR Bit field | Additional MSR   | Event mask value for Front_end_event |
|-------------------------------|------------------------------------|--|--------------------------------------|
| memory_loads                  | None                               | Set TAGLOADS bit in ESCR corresponding to event Uop_Type.      | NBOGUS                               |
| memory_stores                 | None                               | Set TAGSTORES bit in the ESCR corresponding to event Uop_Type. | NBOGUS                               |

**NOTES:**

1. There may be some undercounting of front end events when there is an overflow or underflow of the floating point stack.

**Table 19-35. List of Metrics Available for Execution Tagging (For Execution Event Only)**

| Execution metric              | Upstream ESCR  | TagValue in Upstream ESCR | Event mask value for execution_event |
|-------------------------------|--|---------------------------|--------------------------------------|
| packed_SP_retired             | Set ALL bit in event mask, TagUop bit in ESCR of packed_SP_uop.                | 1                         | NBOGUSO                              |
| packed_DP_retired             | Set ALL bit in event mask, TagUop bit in ESCR of packed_DP_uop.                | 1                         | NBOGUSO                              |
| scalar_SP_retired             | Set ALL bit in event mask, TagUop bit in ESCR of scalar_SP_uop.                | 1                         | NBOGUSO                              |
| scalar_DP_retired             | Set ALL bit in event mask, TagUop bit in ESCR of scalar_DP_uop.                | 1                         | NBOGUSO                              |
| 128_bit_MMX_retired           | Set ALL bit in event mask, TagUop bit in ESCR of 128_bit_MMX_uop.              | 1                         | NBOGUSO                              |
| 64_bit_MMX_retired            | Set ALL bit in event mask, TagUop bit in ESCR of 64_bit_MMX_uop.               | 1                         | NBOGUSO                              |
| X87_FP_retired                | Set ALL bit in event mask, TagUop bit in ESCR of x87_FP_uop.                   | 1                         | NBOGUSO                              |
| X87_SIMD_memory_moves_retired | Set ALLP0, ALLP2 bits in event mask, TagUop bit in ESCR of X87_SIMD_moves_uop. | 1                         | NBOGUSO                              |

**Table 19-36. List of Metrics Available for Replay Tagging (For Replay Event Only)**

| Replay metric <sup>1</sup>                | IA32_PEBS_ENABLE Field to Set  | MSR_PEBS_MATRIX_VERT Bit Field to Set | Additional MSR/ Event | Event Mask Value for Replay_event |
|---|--------------------------------|---------------------------------------|-----------------------|-----------------------------------|
| 1stL_cache_load_miss_retired              | Bit 0, Bit 24, Bit 25          | Bit 0                                 | None                  | NBOGUS                            |
| 2ndL_cache_load_miss_retired <sup>2</sup> | Bit 1, Bit 24, Bit 25          | Bit 0                                 | None                  | NBOGUS                            |
| DTLB_load_miss_retired                    | Bit 2, Bit 24, Bit 25          | Bit 0                                 | None                  | NBOGUS                            |
| DTLB_store_miss_retired                   | Bit 2, Bit 24, Bit 25          | Bit 1                                 | None                  | NBOGUS                            |
| DTLB_all_miss_retired                     | Bit 2, Bit 24, Bit 25          | Bit 0, Bit 1                          | None                  | NBOGUS                            |
| Tagged_mispred_branch                     | Bit 15, Bit 16, Bit 24, Bit 25 | Bit 4                                 | None                  | NBOGUS                            |

**Table 19-36. List of Metrics Available for Replay Tagging (For Replay Event Only) (Contd.)**

| Replay metric <sup>1</sup>           | IA32_PEBS_ENABLE Field to Set | MSR_PEBS_MATRIX_VERT Bit Field to Set | Additional MSR/ Event   | Event Mask Value for Replay_event |
|--------------------------------------|-------------------------------|---------------------------------------|---|-----------------------------------|
| MOB_load_replay_retired <sup>3</sup> | Bit 9, Bit 24, Bit 25         | Bit 0                                 | Select MOB_load_replay event and set PARTIAL_DATA and UNALGN_ADDR bit.                    | NBOGUS                            |
| split_load_retired                   | Bit 10, Bit 24, Bit 25        | Bit 0                                 | Select load_port_replay event with the MSR_SAAT_ESCR1 MSR and set the SPLIT_LD mask bit.  | NBOGUS                            |
| split_store_retired                  | Bit 10, Bit 24, Bit 25        | Bit 1                                 | Select store_port_replay event with the MSR_SAAT_ESCR0 MSR and set the SPLIT_ST mask bit. | NBOGUS                            |

**NOTES:**

1. Certain kinds of  $\mu$ ops cannot be tagged. These include I/O operations, UC and locked accesses, returns, and far transfers.
2. 2nd-level misses retired does not count all 2nd-level misses. It only includes those references that are found to be misses by the fast detection logic and not those that are later found to be misses.
3. While there are several causes for a MOB replay, the event counted with this event mask setting is the case where the data from a load that would otherwise be forwarded is not an aligned subset of the data from a preceding store.

Table 19-37. Event Mask Qualification for Logical Processors

| Event Type     | Event Name          | Event Masks, ESCR[24:9]   | TS or TI   |
|----------------|---------------------|---|--|
| Non-Retirement | BPU_fetch_request   | Bit 0: TCMISS   | TS   |
| Non-Retirement | BSQ_allocation      | Bit<br>0: REQ_TYPE0<br>1: REQ_TYPE1<br>2: REQ_LEN0<br>3: REQ_LEN1<br>5: REQ_IO_TYPE<br>6: REQ_LOCK_TYPE<br>7: REQ_CACHE_TYPE<br>8: REQ_SPLIT_TYPE<br>9: REQ_DEM_TYPE<br>10: REQ_ORD_TYPE<br>11: MEM_TYPE0<br>12: MEM_TYPE1<br>13: MEM_TYPE2 | TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS |
| Non-Retirement | BSQ_cache_reference | Bit<br>0: RD_2ndL_HITS<br>1: RD_2ndL_HITE<br>2: RD_2ndL_HITM<br>3: RD_3rdL_HITS<br>4: RD_3rdL_HITE<br>5: RD_3rdL_HITM<br>6: WR_2ndL_HIT<br>7: WR_3rdL_HIT<br>8: RD_2ndL_MISS<br>9: RD_3rdL_MISS<br>10: WR_2ndL_MISS<br>11: WR_3rdL_MISS     | TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS       |
| Non-Retirement | memory_cancel       | Bit<br>2: ST_RB_FULL<br>3: 64K_CONF   | TS<br>TS   |
| Non-Retirement | SSE_input_assist    | Bit 15: ALL   | TI   |
| Non-Retirement | 64bit_MMX_uop       | Bit 15: ALL   | TI   |
| Non-Retirement | packed_DP_uop       | Bit 15: ALL   | TI   |
| Non-Retirement | packed_SP_uop       | Bit 15: ALL   | TI   |
| Non-Retirement | scalar_DP_uop       | Bit 15: ALL   | TI   |
| Non-Retirement | scalar_SP_uop       | Bit 15: ALL   | TI   |
| Non-Retirement | 128bit_MMX_uop      | Bit 15: ALL   | TI   |
| Non-Retirement | x87_FP_uop          | Bit 15: ALL   | TI   |

**Table 19-37. Event Mask Qualification for Logical Processors (Contd.)**

| Event Type     | Event Name         | Event Masks, ESCR[24:9]   | TS or TI   |
|----------------|--------------------|---|--|
| Non-Retirement | x87_SIMD_moves_uop | Bit<br>3: ALLP0<br>4: ALLP2   | TI<br>TI   |
| Non-Retirement | FSB_data_activity  | Bit<br>0: DRDY_DRV<br>1: DRDY_OWN<br>2: DRDY_OTHER<br>3: DBSY_DRV<br>4: DBSY_OWN<br>5: DBSY_OTHER   | TI<br>TI<br>TI<br>TI<br>TI<br>TI   |
| Non-Retirement | IOQ_allocation     | Bit<br>0: ReqA0<br>1: ReqA1<br>2: ReqA2<br>3: ReqA3<br>4: ReqA4<br>5: ALL_READ<br>6: ALL_WRITE<br>7: MEM_UC<br>8: MEM_WC<br>9: MEM_WT<br>10: MEM_WP<br>11: MEM_WB<br>13: OWN<br>14: OTHER<br>15: PREFETCH | TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS |
| Non-Retirement | IOQ_active_entries | Bit<br>0: ReqA0<br>1: ReqA1<br>2: ReqA2<br>3: ReqA3<br>4: ReqA4<br>5: ALL_READ<br>6: ALL_WRITE<br>7: MEM_UC<br>8: MEM_WC<br>9: MEM_WT<br>10: MEM_WP<br>11: MEM_WB   | TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS<br>TS                         |

Table 19-37. Event Mask Qualification for Logical Processors (Contd.)

| Event Type     | Event Name                  | Event Masks, ESCR[24:9]  | TS or TI                   |
|----------------|-----------------------------|--|----------------------------|
|                |                             | 13: OWN<br>14: OTHER<br>15: PREFETCH   | TS<br>TS<br>TS             |
| Non-Retirement | global_power_events         | Bit 0: RUNNING   | TS                         |
| Non-Retirement | ITLB_reference              | Bit<br>0: HIT<br>1: MISS<br>2: HIT_UC  | TS<br>TS<br>TS             |
| Non-Retirement | MOB_load_replay             | Bit<br>1: NO_STA<br>3: NO_STD<br>4: PARTIAL_DATA<br>5: UNALGN_ADDR               | TS<br>TS<br>TS<br>TS       |
| Non-Retirement | page_walk_type              | Bit<br>0: DTMISS<br>1: ITMISS  | TI<br>TI                   |
| Non-Retirement | uop_type                    | Bit<br>1: TAGLOADS<br>2: TAGSTORES   | TS<br>TS                   |
| Non-Retirement | load_port_replay            | Bit 1: SPLIT_LD  | TS                         |
| Non-Retirement | store_port_replay           | Bit 1: SPLIT_ST  | TS                         |
| Non-Retirement | memory_complete             | Bit<br>0: LSC<br>1: SSC<br>2: USC<br>3: ULC                                      | TS<br>TS<br>TS<br>TS       |
| Non-Retirement | retired_mispred_branch_type | Bit<br>0: UNCONDITIONAL<br>1: CONDITIONAL<br>2: CALL<br>3: RETURN<br>4: INDIRECT | TS<br>TS<br>TS<br>TS<br>TS |
| Non-Retirement | retired_branch_type         | Bit<br>0: UNCONDITIONAL<br>1: CONDITIONAL<br>2: CALL<br>3: RETURN<br>4: INDIRECT | TS<br>TS<br>TS<br>TS<br>TS |

**Table 19-37. Event Mask Qualification for Logical Processors (Contd.)**

| Event Type     | Event Name       | Event Masks, ESCR[24:9]   | TS or TI                                     |
|----------------|------------------|---|--|
| Non-Retirement | tc_ms_xfer       | Bit<br>0: CISC  | TS   |
| Non-Retirement | tc_misc          | Bit<br>4: FLUSH   | TS   |
| Non-Retirement | TC_deliver_mode  | Bit<br>0: DD<br>1: DB<br>2: DI<br>3: BD<br>4: BB<br>5: BI<br>6: ID<br>7: IB | TI<br>TI<br>TI<br>TI<br>TI<br>TI<br>TI<br>TI |
| Non-Retirement | uop_queue_writes | Bit<br>0: FROM_TC_BUILD<br>1: FROM_TC_DELIVER<br>2: FROM_ROM                | TS<br>TS<br>TS                               |
| Non-Retirement | resource_stall   | Bit 5: SBFULL   | TS   |
| Non-Retirement | WC_Buffer        | Bit<br>0: WCB_EVICTS<br>1: WCB_FULL_EVICT<br>2: WCB_HITM_EVICT              | TI<br>TI<br>TI<br>TI                         |
| At Retirement  | instr_retired    | Bit<br>0: NBOGUSNTAG<br>1: NBOGUSTAG<br>2: BOGUSNTAG<br>3: BOGUSTAG         | TS<br>TS<br>TS<br>TS                         |
| At Retirement  | machine_clear    | Bit<br>0: CLEAR<br>2: MOCLEAR<br>6: SMCLEAR                                 | TS<br>TS<br>TS                               |
| At Retirement  | front_end_event  | Bit<br>0: NBOGUS<br>1: BOGUS  | TS<br>TS                                     |
| At Retirement  | replay_event     | Bit<br>0: NBOGUS<br>1: BOGUS  | TS<br>TS                                     |
| At Retirement  | execution_event  | Bit<br>0: NONBOGUS0<br>1: NONBOGUS1   | TS<br>TS                                     |

**Table 19-37. Event Mask Qualification for Logical Processors (Contd.)**

| Event Type    | Event Name             | Event Masks, ESCR[24:9]  | TS or TI                         |
|---------------|------------------------|--|----------------------------------|
|               |                        | 2: NONBOGUS2<br>3: NONBOGUS3<br>4: BOGUS0<br>5: BOGUS1<br>6: BOGUS2<br>7: BOGUS3 | TS<br>TS<br>TS<br>TS<br>TS<br>TS |
| At Retirement | x87_assist             | Bit<br>0: FPSU<br>1: FPSO<br>2: POAO<br>3: POAU<br>4: PREA                       | TS<br>TS<br>TS<br>TS<br>TS       |
| At Retirement | branch_retired         | Bit<br>0: MMNP<br>1: MMNM<br>2: MMTP<br>3: MMTM                                  | TS<br>TS<br>TS<br>TS             |
| At Retirement | mispred_branch_retired | Bit 0: NBOGUS  | TS                               |
| At Retirement | uops_retired           | Bit<br>0: NBOGUS<br>1: BOGUS   | TS<br>TS                         |
| At Retirement | instr_completed        | Bit<br>0: NBOGUS<br>1: BOGUS   | TS<br>TS                         |

## 19.19 PERFORMANCE MONITORING EVENTS FOR INTEL® PENTIUM® M PROCESSORS

The Pentium M processor's performance monitoring events are based on monitoring events for the P6 family of processors. All of these performance events are model specific for the Pentium M processor and are not available in this form in other processors. Table 19-38 lists the performance monitoring events that were added in the Pentium M processor.



**Table 19-38. Performance Monitoring Events on Intel® Pentium® M Processors**

| Name                    | Hex Values | Descriptions   |
|-------------------------|------------|--|
| Power Management        |            |  |
| EMON_EST_TRANS          | 58H        | Number of Enhanced Intel SpeedStep technology transitions:<br>Mask = 00H - All transitions<br>Mask = 02H - Only Frequency transitions      |
| EMON_THERMAL_TRIP       | 59H        | Duration/Occurrences in thermal trip; to count number of thermal trips: bit 22 in PerfEvtSel0/1 needs to be set to enable edge detect.     |
| BPU                     |            |  |
| BR_INST_EXEC            | 88H        | Branch instructions that were executed (not necessarily retired).  |
| BR_MISSP_EXEC           | 89H        | Branch instructions executed that were mispredicted at execution.  |
| BR_BAC_MISSP_EXEC       | 8AH        | Branch instructions executed that were mispredicted at front end (BAC).  |
| BR_CND_EXEC             | 8BH        | Conditional branch instructions that were executed.  |
| BR_CND_MISSP_EXEC       | 8CH        | Conditional branch instructions executed that were mispredicted.   |
| BR_IND_EXEC             | 8DH        | Indirect branch instructions executed.   |
| BR_IND_MISSP_EXEC       | 8EH        | Indirect branch instructions executed that were mispredicted.  |
| BR_RET_EXEC             | 8FH        | Return branch instructions executed.   |
| BR_RET_MISSP_EXEC       | 90H        | Return branch instructions executed that were mispredicted at execution.   |
| BR_RET_BAC_MISSP_EXEC   | 91H        | Return branch instructions executed that were mispredicted at front end (BAC).   |
| BR_CALL_EXEC            | 92H        | CALL instruction executed.   |
| BR_CALL_MISSP_EXEC      | 93H        | CALL instruction executed and miss predicted.  |
| BR_IND_CALL_EXEC        | 94H        | Indirect CALL instructions executed.   |
| Decoder                 |            |  |
| EMON_SIMD_INSTR_RETIRED | CEH        | Number of retired MMX instructions.  |
| EMON_SYNCH_UOPS         | D3H        | Sync micro-ops   |
| EMON_ESP_UOPS           | D7H        | Total number of micro-ops  |
| EMON_FUSED_UOPS_RET     | DAH        | Number of retired fused micro-ops:<br>Mask = 0 - Fused micro-ops<br>Mask = 1 - Only load+Op micro-ops<br>Mask = 2 - Only std+sta micro-ops |
| EMON_UNFUSION           | DBH        | Number of unfusion events in the ROB, happened on a FP exception to a fused $\mu$ op.  |
| Prefetcher              |            |  |
| EMON_PREF_RQSTS_UP      | FOH        | Number of upward prefetches issued.  |
| EMON_PREF_RQSTS_DN      | F8H        | Number of downward prefetches issued.  |

A number of P6 family processor performance monitoring events are modified for the Pentium M processor. Table 19-39 lists the performance monitoring events that were changed in the Pentium M processor, and differ from performance monitoring events for the P6 family of processors.

**Table 19-39. Performance Monitoring Events Modified on Intel® Pentium® M Processors**

| Name                            | Hex Values | Descriptions  |
|---------------------------------|------------|---|
| CPU_CLK_UNHALTED                | 79H        | Number of cycles during which the processor is not halted, and not in a thermal trip.   |
| EMON_SSE_SSE2_INST_RETIRED      | D8H        | Streaming SIMD Extensions Instructions Retired:<br>Mask = 0 - SSE packed single and scalar single<br>Mask = 1 - SSE scalar-single<br>Mask = 2 - SSE2 packed-double<br>Mask = 3 - SSE2 scalar-double   |
| EMON_SSE_SSE2_COMP_INST_RETIRED | D9H        | Computational SSE Instructions Retired:<br>Mask = 0 - SSE packed single<br>Mask = 1 - SSE Scalar-single<br>Mask = 2 - SSE2 packed-double<br>Mask = 3 - SSE2 scalar-double   |
| L2_LD                           | 29H        | L2 data loads   |
| L2_LINES_IN                     | 24H        | L2 lines allocated  |
| L2_LINES_OUT                    | 26H        | L2 lines evicted  |
| L2_M_LINES_OUT                  | 27H        | Lw M-state lines evicted  |
|                                 |            | Mask[0] = 1 - count I state lines<br>Mask[1] = 1 - count S state lines<br>Mask[2] = 1 - count E state lines<br>Mask[3] = 1 - count M state lines<br>Mask[5:4]:<br>00H - Excluding hardware-prefetched lines<br>01H - Hardware-prefetched lines only<br>02H/03H - All (HW-prefetched lines and non HW -- Prefetched lines) |

## 19.20 P6 FAMILY PROCESSOR PERFORMANCE MONITORING EVENTS

Table 19-40 lists the events that can be counted with the performance monitoring counters and read with the RDPMC instruction for the P6 family processors. The unit column gives the microarchitecture or bus unit that produces the event; the event number column gives the hexadecimal number identifying the event; the mnemonic event name column gives the name of the event; the unit mask column gives the unit mask required (if any); the description column describes the event; and the comments column gives additional information about the event.

All of these performance events are model specific for the P6 family processors and are not available in this form in the Pentium 4 processors or the Pentium processors. Some events (such as those added in later generations of the P6 family processors) are only available in specific processors in the P6 family. All performance event encodings not listed in Table 19-40 are reserved and their use will result in undefined counter results.

See the end of the table for notes related to certain entries in the table.

**Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters**

| Unit                         | Event Num. | Mnemonic Event Name  | Unit Mask   | Description  | Comments  |
|------------------------------|------------|----------------------|-------------|--|---|
| Data Cache Unit (DCU)        | 43H        | DATA_MEM_REFS        | 00H         | All loads from any memory type. All stores to any memory type. Each part of a split is counted separately. The internal logic counts not only memory loads and stores, but also internal retries.<br><br>80-bit floating-point accesses are double counted, since they are decomposed into a 16-bit exponent load and a 64-bit mantissa load. Memory accesses are only counted when they are actually performed (such as a load that gets squashed because a previous cache miss is outstanding to the same address, and which finally gets performed, is only counted once).<br><br>Does not include I/O accesses, or other nonmemory accesses. |   |
|                              | 45H        | DCU_LINES_IN         | 00H         | Total lines allocated in DCU.  |   |
|                              | 46H        | DCU_M_LINES_IN       | 00H         | Number of M state lines allocated in DCU.  |   |
|                              | 47H        | DCU_M_LINES_OUT      | 00H         | Number of M state lines evicted from DCU. This includes evictions via snoop HITM, intervention or replacement.   |   |
|                              | 48H        | DCU_MISS_OUTSTANDING | 00H         | Weighted number of cycles while a DCU miss is outstanding, incremented by the number of outstanding cache misses at any particular time.<br><br>Cacheable read requests only are considered.<br><br>Uncacheable requests are excluded.<br><br>Read-for-ownerships are counted, as well as line fills, invalidates, and stores.   | An access that also misses the L2 is short-changed by 2 cycles (i.e., if counts N cycles, should be N+2 cycles).<br><br>Subsequent loads to the same cache line will not result in any additional counts.<br><br>Count value not precise, but still useful. |
| Instruction Fetch Unit (IFU) | 80H        | IFU_IFETCH           | 00H         | Number of instruction fetches, both cacheable and noncacheable, including UC fetches.  |   |
|                              | 81H        | IFU_IFETCH_MISS      | 00H         | Number of instruction fetch misses<br>All instruction fetches that do not hit the IFU (i.e., that produce memory requests). This includes UC accesses.   |   |
|                              | 85H        | ITLB_MISS            | 00H         | Number of ITLB misses.   |   |
|                              | 86H        | IFU_MEM_STALL        | 00H         | Number of cycles instruction fetch is stalled, for any reason.<br><br>Includes IFU cache misses, ITLB misses, ITLB faults, and other minor stalls.   |   |
|                              | 87H        | ILD_STALL            | 00H         | Number of cycles that the instruction length decoder is stalled.   |   |
| L2 Cache <sup>1</sup>        | 28H        | L2_IFETCH            | MESI<br>0FH | Number of L2 instruction fetches.<br><br>This event indicates that a normal instruction fetch was received by the L2.  |   |

Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)

| Unit                                  | Event Num. | Mnemonic Event Name | Unit Mask               | Description  | Comments  |
|---------------------------------------|------------|---------------------|-------------------------|--|---|
|                                       |            |                     |                         | The count includes only L2 cacheable instruction fetches; it does not include UC instruction fetches.<br>It does not include ITLB miss accesses.   |   |
|                                       | 29H        | L2_LD               | MESI<br>0FH             | Number of L2 data loads.<br>This event indicates that a normal, unlocked, load memory access was received by the L2.<br>It includes only L2 cacheable memory accesses; it does not include I/O accesses, other nonmemory accesses, or memory accesses such as UC/WT memory accesses.<br>It does include L2 cacheable TLB miss memory accesses.   |   |
|                                       | 2AH        | L2_ST               | MESI<br>0FH             | Number of L2 data stores.<br>This event indicates that a normal, unlocked, store memory access was received by the L2.<br><br>it indicates that the DCU sent a read-for-ownership request to the L2. It also includes Invalid to Modified requests sent by the DCU to the L2.<br><br>It includes only L2 cacheable memory accesses; it does not include I/O accesses, other nonmemory accesses, or memory accesses such as UC/WT memory accesses.<br>It includes TLB miss memory accesses. |   |
|                                       | 24H        | L2_LINES_IN         | 00H                     | Number of lines allocated in the L2.   |   |
|                                       | 26H        | L2_LINES_OUT        | 00H                     | Number of lines removed from the L2 for any reason.  |   |
|                                       | 25H        | L2_M_LINES_INM      | 00H                     | Number of modified lines allocated in the L2.  |   |
|                                       | 27H        | L2_M_LINES_OUTM     | 00H                     | Number of modified lines removed from the L2 for any reason.   |   |
|                                       | 2EH        | L2_RQSTS            | MESI<br>0FH             | Total number of L2 requests.   |   |
|                                       | 21H        | L2_ADS              | 00H                     | Number of L2 address strobes.  |   |
|                                       | 22H        | L2_DBUS_BUSY        | 00H                     | Number of cycles during which the L2 cache data bus was busy.  |   |
|                                       | 23H        | L2_DBUS_BUSY_RD     | 00H                     | Number of cycles during which the data bus was busy transferring read data from L2 to the processor.   |   |
| External Bus Logic (EBL) <sup>2</sup> | 62H        | BUS_DRDY_CLOCKS     | 00H (Self)<br>20H (Any) | Number of clocks during which DRDY# is asserted.<br>Utilization of the external system data bus during data transfers.   | Unit Mask = 00H counts bus clocks when the processor is driving DRDY#.<br>Unit Mask = 20H counts in processor clocks when any agent is driving DRDY#. |

**Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)**

| Unit | Event Num. | Mnemonic Event Name | Unit Mask               | Description  | Comments   |
|------|------------|---------------------|-------------------------|--|--|
|      | 63H        | BUS_LOCK_CLOCKS     | 00H (Self)<br>20H (Any) | Number of clocks during which LOCK# is asserted on the external system bus. <sup>3</sup>   | Always counts in processor clocks.   |
|      | 60H        | BUS_REQ_OUTSTANDING | 00H (Self)              | Number of bus requests outstanding. This counter is incremented by the number of cacheable read bus requests outstanding in any given cycle. | Counts only DCU full-line cacheable reads, not RFOs, writes, instruction fetches, or anything else. Counts "waiting for bus to complete" (last data chunk received). |
|      | 65H        | BUS_TRAN_BRD        | 00H (Self)<br>20H (Any) | Number of burst read transactions.   |  |
|      | 66H        | BUS_TRAN_RFO        | 00H (Self)<br>20H (Any) | Number of completed read for ownership transactions.   |  |
|      | 67H        | BUS_TRANS_WB        | 00H (Self)<br>20H (Any) | Number of completed write back transactions.   |  |
|      | 68H        | BUS_TRAN_IFETCH     | 00H (Self)<br>20H (Any) | Number of completed instruction fetch transactions.  |  |
|      | 69H        | BUS_TRAN_INVALID    | 00H (Self)<br>20H (Any) | Number of completed invalidate transactions.   |  |
|      | 6AH        | BUS_TRAN_PWR        | 00H (Self)<br>20H (Any) | Number of completed partial write transactions.  |  |
|      | 6BH        | BUS_TRANS_P         | 00H (Self)<br>20H (Any) | Number of completed partial transactions.  |  |
|      | 6CH        | BUS_TRANS_IO        | 00H (Self)<br>20H (Any) | Number of completed I/O transactions.  |  |
|      | 6DH        | BUS_TRAN_DEF        | 00H (Self)<br>20H (Any) | Number of completed deferred transactions.   |  |

Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)

| Unit | Event Num. | Mnemonic Event Name | Unit Mask               | Description  | Comments   |
|------|------------|---------------------|-------------------------|--|--|
|      | 6EH        | BUS_TRAN_BURST      | 00H (Self)<br>20H (Any) | Number of completed burst transactions.  |  |
|      | 70H        | BUS_TRAN_ANY        | 00H (Self)<br>20H (Any) | Number of all completed bus transactions.<br>Address bus utilization can be calculated knowing the minimum address bus occupancy.<br>Includes special cycles, etc. |  |
|      | 6FH        | BUS_TRAN_MEM        | 00H (Self)<br>20H (Any) | Number of completed memory transactions.   |  |
|      | 64H        | BUS_DATA_RCV        | 00H (Self)              | Number of bus clock cycles during which this processor is receiving data.  |  |
|      | 61H        | BUS_BNR_DRV         | 00H (Self)              | Number of bus clock cycles during which this processor is driving the BNR# pin.  |  |
|      | 7AH        | BUS_HIT_DRV         | 00H (Self)              | Number of bus clock cycles during which this processor is driving the HIT# pin.  | Includes cycles due to snoop stalls.<br>The event counts correctly, but BPM <sub>i</sub> (breakpoint monitor) pins function as follows based on the setting of the PC bits (bit 19 in the PerfEvtSel0 and PerfEvtSel1 registers): <ul style="list-style-type: none"> <li>▪ If the core-clock-to-bus-clock ratio is 2:1 or 3:1, and a PC bit is set, the BPM<sub>i</sub> pins will be asserted for a single clock when the counters overflow.</li> <li>▪ If the PC bit is clear, the processor toggles the BPM<sub>i</sub> pins when the counter overflows.</li> <li>▪ If the clock ratio is not 2:1 or 3:1, the BPM<sub>i</sub> pins will not function for these performance monitoring counter events.</li> </ul> |
|      | 7BH        | BUS_HITM_DRV        | 00H (Self)              | Number of bus clock cycles during which this processor is driving the HITM# pin.   | Includes cycles due to snoop stalls.<br>The event counts correctly, but BPM <sub>i</sub> (breakpoint monitor) pins function as follows based on the setting of the PC bits (bit 19 in the PerfEvtSel0 and PerfEvtSel1 registers): <ul style="list-style-type: none"> <li>▪ If the core-clock-to-bus-clock ratio is 2:1 or 3:1, and a PC bit is set, the BPM<sub>i</sub> pins will be asserted for a single clock when the counters overflow.</li> </ul>  |

**Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)**

| Unit                | Event Num. | Mnemonic Event Name | Unit Mask  | Description  | Comments  |
|---------------------|------------|---------------------|------------|--|---|
|                     |            |                     |            |  | <ul style="list-style-type: none"> <li>If the PC bit is clear, the processor toggles the BPM pins when the counter overflows.</li> <li>If the clock ratio is not 2:1 or 3:1, the BPM pins will not function for these performance monitoring counter events.</li> </ul> |
|                     | 7EH        | BUS_SNOOP_STALL     | 00H (Self) | Number of clock cycles during which the bus is snoop stalled.  |   |
| Floating-Point Unit | C1H        | FLOPS               | 00H        | <p>Number of computational floating-point operations retired.</p> <p>Excludes floating-point computational operations that cause traps or assists.</p> <p>Includes floating-point computational operations executed by the assist handler.</p> <p>Includes internal sub-operations for complex floating-point instructions like transcendentals.</p> <p>Excludes floating-point loads and stores.</p>        | Counter 0 only.   |
|                     | 10H        | FP_COMP_OPS_EXE     | 00H        | <p>Number of computational floating-point operations executed.</p> <p>The number of FADD, FSUB, FCOM, FMULs, integer MULs and IMULs, FDIVs, FPREM, FSQRTS, integer DIVs, and IDIVs.</p> <p>This number does not include the number of cycles, but the number of operations.</p> <p>This event does not distinguish an FADD used in the middle of a transcendental flow from a separate FADD instruction.</p> | Counter 0 only.   |
|                     | 11H        | FP_ASSIST           | 00H        | Number of floating-point exception cases handled by microcode.   | Counter 1 only.<br>This event includes counts due to speculative execution.   |
|                     | 12H        | MUL                 | 00H        | <p>Number of multiplies.</p> <p>This count includes integer as well as FP multiplies and is speculative.</p>   | Counter 1 only.   |
|                     | 13H        | DIV                 | 00H        | <p>Number of divides.</p> <p>This count includes integer as well as FP divides and is speculative.</p>   | Counter 1 only.   |
|                     | 14H        | CYCLES_DIV_BUSY     | 00H        | <p>Number of cycles during which the divider is busy, and cannot accept new divides.</p> <p>This includes integer and FP divides, FPREM, FPSQRT, etc. and is speculative.</p>  | Counter 0 only.   |

Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)

| Unit                                | Event Num. | Mnemonic Event Name      | Unit Mask                | Description   | Comments   |
|-------------------------------------|------------|--------------------------|--------------------------|---|--|
| Memory Ordering                     | 03H        | LD_BLOCKS                | 00H                      | Number of load operations delayed due to store buffer blocks.<br><br>Includes counts caused by preceding stores whose addresses are unknown, preceding stores whose addresses are known but whose data is unknown, and preceding stores that conflicts with the load but which incompletely overlap the load.   |  |
|                                     | 04H        | SB_DRAINS                | 00H                      | Number of store buffer drain cycles.<br><br>Incremented every cycle the store buffer is draining.<br><br>Draining is caused by serializing operations like CPUID, synchronizing operations like XCHG, interrupt acknowledgment, as well as other conditions (such as cache flushing).   |  |
|                                     | 05H        | MISALIGN_MEM_REF         | 00H                      | Number of misaligned data memory references.<br><br>Incremented by 1 every cycle, during which either the processor's load or store pipeline dispatches a misaligned $\mu$ op.<br><br>Counting is performed if it is the first or second half, or if it is blocked, squashed, or missed.<br><br>In this context, misaligned means crossing a 64-bit boundary. | MISALIGN_MEM_REF is only an approximation to the true number of misaligned memory references.<br><br>The value returned is roughly proportional to the number of misaligned memory accesses (the size of the problem). |
|                                     | 07H        | EMON_KNI_PREF_DISPATCHED | 00H<br>01H<br>02H<br>03H | Number of Streaming SIMD extensions prefetch/weakly-ordered instructions dispatched (speculative prefetches are included in counting):<br><br>0: prefetch NTA<br>1: prefetch T1<br>2: prefetch T2<br>3: weakly ordered stores   | Counters 0 and 1. Pentium III processor only.  |
|                                     | 4BH        | EMON_KNI_PREF_MISS       | 00H<br>01H<br>02H<br>03H | Number of prefetch/weakly-ordered instructions that miss all caches:<br><br>0: prefetch NTA<br>1: prefetch T1<br>2: prefetch T2<br>3: weakly ordered stores   | Counters 0 and 1. Pentium III processor only.  |
| Instruction Decoding and Retirement | COH        | INST_RETIRED             | 00H                      | Number of instructions retired.   | A hardware interrupt received during/after the last iteration of the REP STOS flow causes the counter to undercount by 1 instruction.  |
|                                     |            |                          |                          |   | An SMI received while executing a HLT instruction will cause the performance counter to not count the RSM instruction and undercount by 1.   |



**Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)**

| Unit       | Event Num. | Mnemonic Event Name           | Unit Mask  | Description  | Comments                                      |
|------------|------------|-------------------------------|------------|--|---|
|            | C2H        | UOPS_RETIRE                   | 00H        | Number of $\mu$ ops retired.   |   |
|            | D0H        | INST_DECODED                  | 00H        | Number of instructions decoded.  |   |
|            | D8H        | EMON_KNI_INST_RETIRE          | 00H<br>01H | Number of Streaming SIMD extensions retired:<br>0: packed & scalar<br>1: scalar  | Counters 0 and 1. Pentium III processor only. |
|            | D9H        | EMON_KNI_COMP_INST_RET        | 00H<br>01H | Number of Streaming SIMD extensions computation instructions retired:<br>0: packed and scalar<br>1: scalar   | Counters 0 and 1. Pentium III processor only. |
| Interrupts | C8H        | HW_INT_RX                     | 00H        | Number of hardware interrupts received.  |   |
|            | C6H        | CYCLES_INT_MASKED             | 00H        | Number of processor cycles for which interrupts are disabled.  |   |
|            | C7H        | CYCLES_INT_PENDING_AND_MASKED | 00H        | Number of processor cycles for which interrupts are disabled and interrupts are pending.   |   |
| Branches   | C4H        | BR_INST_RETIRE                | 00H        | Number of branch instructions retired.   |   |
|            | C5H        | BR_MISS_PRED_RETIRE           | 00H        | Number of mispredicted branches retired.   |   |
|            | C9H        | BR_TAKEN_RETIRE               | 00H        | Number of taken branches retired.  |   |
|            | CAH        | BR_MISS_PRED_TAKEN_RET        | 00H        | Number of taken mispredictions branches retired.   |   |
|            | E0H        | BR_INST_DECODED               | 00H        | Number of branch instructions decoded.   |   |
|            | E2H        | BTB_MISSES                    | 00H        | Number of branches for which the BTB did not produce a prediction.   |   |
|            | E4H        | BR_BOGUS                      | 00H        | Number of bogus branches.  |   |
|            | E6H        | BACLEAR                       | 00H        | Number of times BACLEAR is asserted.<br>This is the number of times that a static branch prediction was made, in which the branch decoder decided to make a branch prediction because the BTB did not. |   |
| Stalls     | A2H        | RESOURCE_STALLS               | 00H        | Incremented by 1 during every cycle for which there is a resource related stall.<br>Includes register renaming buffer entries, memory buffer entries.  |   |

Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)

| Unit                      | Event Num.  | Mnemonic Event Name                          | Unit Mask  | Description  | Comments  |
|---------------------------|---|--|--|--|---|
|                           |   |  |  | Does not include stalls due to bus queue full, too many cache misses, etc.<br>In addition to resource related stalls, this event counts some other events.<br>Includes stalls arising during branch misprediction recovery, such as if retirement of the mispredicted branch is delayed and stalls arising while store buffer is draining from synchronizing operations. |   |
|                           | D2H   | PARTIAL_RAT_STALLS                           | 00H  | Number of cycles or events for partial stalls. This includes flag partial stalls.  |   |
| Segment Register Loads    | 06H   | SEGMENT_REG_LOADS                            | 00H  | Number of segment register loads.  |   |
| Clocks                    | 79H   | CPU_CLK_UNHALTED                             | 00H  | Number of cycles during which the processor is not halted.   |   |
| MMX Unit                  | B0H   | MMX_INSTR_EXEC                               | 00H  | Number of MMX Instructions Executed.   | Available in Intel Celeron, Pentium II and Pentium II Xeon processors only.<br>Does not account for MOVQ and MOVD stores from register to memory. |
|                           | B1H   | MMX_SAT_INSTR_EXEC                           | 00H  | Number of MMX Saturating Instructions Executed.  | Available in Pentium II and Pentium III processors only.  |
|                           | B2H   | MMX_UOPS_EXEC                                | 0FH  | Number of MMX $\mu$ ops Executed.  | Available in Pentium II and Pentium III processors only.  |
|                           | B3H   | MMX_INSTR_TYPE_EXEC                          | 01H  | MMX packed multiply instructions executed.   | Available in Pentium II and Pentium III processors only.  |
|                           |   |  | 02H  | MMX packed shift instructions executed.  |   |
|                           |   |  | 04H  | MMX pack operation instructions executed.  |   |
|                           |   |  | 08H  | MMX unpack operation instructions executed.  |   |
|                           |   |  | 10H  | MMX packed logical instructions executed.  |   |
|                           | 20H   | MMX packed arithmetic instructions executed. |  |  |   |
| CCH                       | FP_MMX_TRANS  | 00H  | Transitions from MMX instruction to floating-point instructions.           | Available in Pentium II and Pentium III processors only.   |   |
| 01H                       | Transitions from floating-point instructions to MMX instructions. |  |  |  |   |
| CDH                       | MMX_ASSIST  | 00H  | Number of MMX Assists (that is, the number of EMMS instructions executed). | Available in Pentium II and Pentium III processors only.   |   |
| CEH                       | MMX_INSTR_RET   | 00H  | Number of MMX Instructions Retired.  | Available in Pentium II processors only.   |   |
| Segment Register Renaming | D4H   | SEG_RENAME_STALLS                            |  | Number of Segment Register Renaming Stalls:  | Available in Pentium II and Pentium III processors only.  |

**Table 19-40. Events That Can Be Counted with the P6 Family Performance Monitoring Counters (Contd.)**

| Unit | Event Num. | Mnemonic Event Name | Unit Mask                       | Description   | Comments   |
|------|------------|---------------------|---------------------------------|---|--|
|      |            |                     | 02H<br>04H<br>08H<br>0FH        | Segment register ES<br>Segment register DS<br>Segment register FS<br>Segment register FS<br>Segment registers<br>ES + DS + FS + GS  |  |
|      | D5H        | SEG_REG_RENAMES     | 01H<br>02H<br>04H<br>08H<br>0FH | Number of Segment Register Renames:<br>Segment register ES<br>Segment register DS<br>Segment register FS<br>Segment register FS<br>Segment registers<br>ES + DS + FS + GS | Available in Pentium II and Pentium III processors only. |
|      | D6H        | RET_SEG_RENAMES     | 00H                             | Number of segment register rename events retired.   | Available in Pentium II and Pentium III processors only. |

**NOTES:**

- Several L2 cache events, where noted, can be further qualified using the Unit Mask (UMSK) field in the PerfEvtSel0 and PerfEvtSel1 registers. The lower 4 bits of the Unit Mask field are used in conjunction with L2 events to indicate the cache state or cache states involved.  
The P6 family processors identify cache states using the "MESI" protocol and consequently each bit in the Unit Mask field represents one of the four states: UMSK[3] = M (8H) state, UMSK[2] = E (4H) state, UMSK[1] = S (2H) state, and UMSK[0] = I (1H) state. UMSK[3:0] = MESI" (FH) should be used to collect data for all states; UMSK = 0H, for the applicable events, will result in nothing being counted.
- All of the external bus logic (EBL) events, except where noted, can be further qualified using the Unit Mask (UMSK) field in the PerfEvtSel0 and PerfEvtSel1 registers.  
Bit 5 of the UMSK field is used in conjunction with the EBL events to indicate whether the processor should count transactions that are self-generated (UMSK[5] = 0) or transactions that result from any processor on the bus (UMSK[5] = 1).
- L2 cache locks, so it is possible to have a zero count.

## 19.21 PENTIUM PROCESSOR PERFORMANCE MONITORING EVENTS

Table 19-41 lists the events that can be counted with the performance monitoring counters for the Pentium processor. The Event Number column gives the hexadecimal code that identifies the event and that is entered in the ES0 or ES1 (event select) fields of the CESR MSR. The Mnemonic Event Name column gives the name of the event, and the Description and Comments columns give detailed descriptions of the events. Most events can be counted with either counter 0 or counter 1; however, some events can only be counted with only counter 0 or only counter 1 (as noted).

**NOTE**

The events in the table that are shaded are implemented only in the Pentium processor with MMX technology.

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters**

| Event Num. | Mnemonic Event Name                      | Description   | Comments  |
|------------|--|---|---|
| 00H        | DATA_READ                                | Number of memory data reads (internal data cache hit and miss combined).  | Split cycle reads are counted individually. Data Memory Reads that are part of TLB miss processing are not included. These events may occur at a maximum of two per clock. I/O is not included.   |
| 01H        | DATA_WRITE                               | Number of memory data writes (internal data cache hit and miss combined); I/O not included.                               | Split cycle writes are counted individually. These events may occur at a maximum of two per clock. I/O is not included.   |
| 0H2        | DATA_TLB_MISS                            | Number of misses to the data cache translation look-aside buffer.   |   |
| 03H        | DATA_READ_MISS                           | Number of memory read accesses that miss the internal data cache whether or not the access is cacheable or noncacheable.  | Additional reads to the same cache line after the first BRDY# of the burst line fill is returned but before the final (fourth) BRDY# has been returned, will not cause the counter to be incremented additional times.<br>Data accesses that are part of TLB miss processing are not included. Accesses directed to I/O space are not included. |
| 04H        | DATA WRITE MISS                          | Number of memory write accesses that miss the internal data cache whether or not the access is cacheable or noncacheable. | Data accesses that are part of TLB miss processing are not included. Accesses directed to I/O space are not included.   |
| 05H        | WRITE_HIT_TO_M_OR_E-STATE_LINES          | Number of write hits to exclusive or modified lines in the data cache.  | These are the writes that may be held up if EWBE# is inactive. These events may occur a maximum of two per clock.   |
| 06H        | DATA_CACHE_LINES_WRITTEN_BACK            | Number of dirty lines (all) that are written back, regardless of the cause.   | Replacements and internal and external snoops can all cause writeback and are counted.  |
| 07H        | EXTERNAL_SNOOPS                          | Number of accepted external snoops whether they hit in the code cache or data cache or neither.                           | Assertions of EADS# outside of the sampling interval are not counted, and no internal snoops are counted.   |
| 08H        | EXTERNAL_DATA_CACHE_SNOOP_HITS           | Number of external snoops to the data cache.  | Snoop hits to a valid line in either the data cache, the data line fill buffer, or one of the write back buffers are all counted as hits.   |
| 09H        | MEMORY ACCESSES IN BOTH PIPES            | Number of data memory reads or writes that are paired in both pipes of the pipeline.                                      | These accesses are not necessarily run in parallel due to cache misses, bank conflicts, etc.  |
| 0AH        | BANK CONFLICTS                           | Number of actual bank conflicts.  |   |
| 0BH        | MISALIGNED DATA MEMORY OR I/O REFERENCES | Number of memory or I/O reads or writes that are misaligned.  | A 2- or 4-byte access is misaligned when it crosses a 4-byte boundary; an 8-byte access is misaligned when it crosses an 8-byte boundary. Ten byte accesses are treated as two separate accesses of 8 and 2 bytes each.   |
| 0CH        | CODE READ                                | Number of instruction reads; whether the read is cacheable or noncacheable.   | Individual 8-byte noncacheable instruction reads are counted.   |
| 0DH        | CODE TLB MISS                            | Number of instruction reads that miss the code TLB whether the read is cacheable or noncacheable.                         | Individual 8-byte noncacheable instruction reads are counted.   |
| 0EH        | CODE CACHE MISS                          | Number of instruction reads that miss the internal code cache; whether the read is cacheable or noncacheable.             | Individual 8-byte noncacheable instruction reads are counted.   |

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name              | Description  | Comments   |
|------------|----------------------------------|--|--|
| 0FH        | ANY SEGMENT REGISTER LOADED      | Number of writes into any segment register in real or protected mode including the LDTR, GDTR, IDTR, and TR.   | Segment loads are caused by explicit segment register load instructions, far control transfers, and task switches. Far control transfers and task switches causing a privilege level change will signal this event twice. Interrupts and exceptions may initiate a far control transfer.   |
| 10H        | Reserved                         |  |  |
| 11H        | Reserved                         |  |  |
| 12H        | Branches                         | Number of taken and not taken branches, including: conditional branches, jumps, calls, returns, software interrupts, and interrupt returns.                                      | Also counted as taken branches are serializing instructions, VERR and VERW instructions, some segment descriptor loads, hardware interrupts (including FLUSH#), and programmatic exceptions that invoke a trap or fault handler. The pipe is not necessarily flushed.<br>The number of branches actually executed is measured, not the number of predicted branches.   |
| 13H        | BTB_HITS                         | Number of BTB hits that occur.   | Hits are counted only for those instructions that are actually executed.   |
| 14H        | TAKEN_BRANCH_OR_BTBT_HIT         | Number of taken branches or BTB hits that occur.   | This event type is a logical OR of taken branches and BTB hits. It represents an event that may cause a hit in the BTB. Specifically, it is either a candidate for a space in the BTB or it is already in the BTB.   |
| 15H        | PIPELINE FLUSHES                 | Number of pipeline flushes that occur<br>Pipeline flushes are caused by BTB misses on taken branches, mispredictions, exceptions, interrupts, and some segment descriptor loads. | The counter will not be incremented for serializing instructions (serializing instructions cause the prefetch queue to be flushed but will not trigger the Pipeline Flushed event counter) and software interrupts (software interrupts do not flush the pipeline).  |
| 16H        | INSTRUCTIONS_EXECUTED            | Number of instructions executed (up to two per clock).   | Invocations of a fault handler are considered instructions. All hardware and software interrupts and exceptions will also cause the count to be incremented. Repeat prefixed string instructions will only increment this counter once despite the fact that the repeat loop executes the same instruction multiple times until the loop criteria is satisfied.<br>This applies to all the Repeat string instruction prefixes (i.e., REP, REPE, REPZ, REPNE, and REPNZ). This counter will also only increment once per each HLT instruction executed regardless of how many cycles the processor remains in the HALT state. |
| 17H        | INSTRUCTIONS_EXECUTED_V PIPE     | Number of instructions executed in the V_pipe.<br>The event indicates the number of instructions that were paired.   | This event is the same as the 16H event except it only counts the number of instructions actually executed in the V-pipe.  |
| 18H        | BUS_CYCLE_DURATION               | Number of clocks while a bus cycle is in progress.<br>This event measures bus use.   | The count includes HLDA, AHOLD, and BOFF# clocks.  |
| 19H        | WRITE_BUFFER_FULL_STALL_DURATION | Number of clocks while the pipeline is stalled due to full write buffers.  | Full write buffers stall data memory read misses, data memory write misses, and data memory write hits to S-state lines. Stalls on I/O accesses are not included.  |

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name                         | Description   | Comments   |
|------------|---|---|--|
| 1AH        | WAITING_FOR_DATA_MEMORY_READ_STALL_DURATION | Number of clocks while the pipeline is stalled while waiting for data memory reads.   | Data TLB Miss processing is also included in the count. The pipeline stalls while a data memory read is in progress including attempts to read that are not bypassed while a line is being filled.   |
| 1BH        | STALL ON WRITE TO AN E- OR M-STATE LINE     | Number of stalls on writes to E- or M-state lines.  |  |
| 1CH        | LOCKED BUS CYCLE                            | Number of locked bus cycles that occur as the result of the LOCK prefix or LOCK instruction, page-table updates, and descriptor table updates.                            | Only the read portion of the locked read-modify-write is counted. Split locked cycles (SCYC active) count as two separate accesses. Cycles restarted due to BOFF# are not re-counted.  |
| 1DH        | I/O READ OR WRITE CYCLE                     | Number of bus cycles directed to I/O space.   | Misaligned I/O accesses will generate two bus cycles. Bus cycles restarted due to BOFF# are not re-counted.  |
| 1EH        | NONCACHEABLE_MEMORY_READS                   | Number of noncacheable instruction or data memory read bus cycles.<br>The count includes read cycles caused by TLB misses, but does not include read cycles to I/O space. | Cycles restarted due to BOFF# are not re-counted.  |
| 1FH        | PIPELINE_AGI_STALLS                         | Number of address generation interlock (AGI) stalls.<br>An AGI occurring in both the U- and V-pipelines in the same clock signals this event twice.                       | An AGI occurs when the instruction in the execute stage of either of U- or V-pipelines is writing to either the index or base address register of an instruction in the D2 (address generation) stage of either the U- or V- pipelines.  |
| 20H        | Reserved                                    |   |  |
| 21H        | Reserved                                    |   |  |
| 22H        | FLOPS                                       | Number of floating-point operations that occur.   | Number of floating-point adds, subtracts, multiplies, divides, remainders, and square roots are counted. The transcendental instructions consist of multiple adds and multiplies and will signal this event multiple times. Instructions generating the divide-by-zero, negative square root, special operand, or stack exceptions will not be counted.<br><br>Instructions generating all other floating-point exceptions will be counted. The integer multiply instructions and other instructions which use the x87 FPU will be counted.                                    |
| 23H        | BREAKPOINT MATCH ON DRO REGISTER            | Number of matches on register DRO breakpoint.   | The counters is incremented regardless if the breakpoints are enabled or not. However, if breakpoints are not enabled, code breakpoint matches will not be checked for instructions executed in the V-pipe and will not cause this counter to be incremented. (They are checked on instruction executed in the U-pipe only when breakpoints are not enabled.)<br><br>These events correspond to the signals driven on the BP[3:0] pins. Refer to Chapter 17, "Debug, Branch Profile, TSC, and Intel® Resource Director Technology (Intel® RDT) Features" for more information. |
| 24H        | BREAKPOINT MATCH ON DR1 REGISTER            | Number of matches on register DR1 breakpoint.   | See comment for 23H event.   |

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name                          | Description  | Comments  |
|------------|--|--|---|
| 25H        | BREAKPOINT MATCH ON DR2 REGISTER             | Number of matches on register DR2 breakpoint.  | See comment for 23H event.  |
| 26H        | BREAKPOINT MATCH ON DR3 REGISTER             | Number of matches on register DR3 breakpoint.  | See comment for 23H event.  |
| 27H        | HARDWARE INTERRUPTS                          | Number of taken INTR and NMI interrupts.   |   |
| 28H        | DATA_READ_OR_WRITE                           | Number of memory data reads and/or writes (internal data cache hit and miss combined).   | Split cycle reads and writes are counted individually. Data Memory Reads that are part of TLB miss processing are not included. These events may occur at a maximum of two per clock. I/O is not included.  |
| 29H        | DATA_READ_MISS OR_WRITE MISS                 | Number of memory read and/or write accesses that miss the internal data cache, whether or not the access is cacheable or noncacheable.                               | Additional reads to the same cache line after the first BRDY# of the burst line fill is returned but before the final (fourth) BRDY# has been returned, will not cause the counter to be incremented additional times.<br>Data accesses that are part of TLB miss processing are not included. Accesses directed to I/O space are not included. |
| 2AH        | BUS_OWNERSHIP_LATENCY (Counter 0)            | The time from LRM bus ownership request to bus ownership granted (that is, the time from the earlier of a PBREQ (0), PHITM# or HITM# assertion to a PBGNT assertion) | The ratio of the 2AH events counted on counter 0 and counter 1 is the average stall time due to bus ownership conflict.   |
| 2AH        | BUS OWNERSHIP TRANSFERS (Counter 1)          | The number of buss ownership transfers (that is, the number of PBREQ (0) assertions  | The ratio of the 2AH events counted on counter 0 and counter 1 is the average stall time due to bus ownership conflict.   |
| 2BH        | MMX_INSTRUCTIONS_EXECUTED_U-PIPE (Counter 0) | Number of MMX instructions executed in the U-pipe  |   |
| 2BH        | MMX_INSTRUCTIONS_EXECUTED_V-PIPE (Counter 1) | Number of MMX instructions executed in the V-pipe  |   |
| 2CH        | CACHE_M-STATE_LINE_SHARING (Counter 0)       | Number of times a processor identified a hit to a modified line due to a memory access in the other processor (PHITM (0))  | If the average memory latencies of the system are known, this event enables the user to count the Write Backs on PHITM(0) penalty and the Latency on Hit Modified(I) penalty.   |
| 2CH        | CACHE_LINE_SHARING (Counter 1)               | Number of shared data lines in the L1 cache (PHIT (0))   |   |
| 2DH        | EMMS_INSTRUCTIONS_EXECUTED (Counter 0)       | Number of EMMS instructions executed   |   |

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name  | Description   | Comments   |
|------------|--|---|--|
| 2DH        | TRANSITIONS_BETWEEN_MMX_AND_FP_INSTRUCTIONS<br>(Counter 1) | Number of transitions between MMX and floating-point instructions or vice versa<br><br>An even count indicates the processor is in MMX state. an odd count indicates it is in FP state. | This event counts the first floating-point instruction following an MMX instruction or first MMX instruction following a floating-point instruction.<br><br>The count may be used to estimate the penalty in transitions between floating-point state and MMX state.   |
| 2EH        | BUS_UTILIZATION_DUE_TO_PROCESSOR_ACTIVITY<br>(Counter 0)   | Number of clocks the bus is busy due to the processor's own activity (the bus activity that is caused by the processor)   |  |
| 2EH        | WRITES_TO_NONCACHEABLE_MEMORY<br>(Counter 1)               | Number of write accesses to noncacheable memory   | The count includes write cycles caused by TLB misses and I/O write cycles.<br><br>Cycles restarted due to BOFF# are not re-counted.  |
| 2FH        | SATURATING_MMX_INSTRUCTIONS_EXECUTED<br>(Counter 0)        | Number of saturating MMX instructions executed, independently of whether they actually saturated.   |  |
| 2FH        | SATURATIONS_PERFORMED<br>(Counter 1)                       | Number of MMX instructions that used saturating arithmetic when at least one of its results actually saturated  | If an MMX instruction operating on 4 doublewords saturated in three out of the four results, the counter will be incremented by one only.  |
| 30H        | NUMBER_OF_CYCLES_NOT_IN_HALT_STATE<br>(Counter 0)          | Number of cycles the processor is not idle due to HLT instruction   | This event will enable the user to calculate "net CPI". Note that during the time that the processor is executing the HLT instruction, the Time-Stamp Counter is not disabled. Since this event is controlled by the Counter Controls CCO, CC1 it can be used to calculate the CPI at CPL=3, which the TSC cannot provide. |
| 30H        | DATA_CACHE_TLB_MISS_STALL_DURATION<br>(Counter 1)          | Number of clocks the pipeline is stalled due to a data cache translation look-aside buffer (TLB) miss   |  |
| 31H        | MMX_INSTRUCTION_DATA_READS<br>(Counter 0)                  | Number of MMX instruction data reads  |  |
| 31H        | MMX_INSTRUCTION_DATA_READ_MISSES<br>(Counter 1)            | Number of MMX instruction data read misses  |  |
| 32H        | FLOATING_POINT_STALLS_DURATION<br>(Counter 0)              | Number of clocks while pipe is stalled due to a floating-point freeze   |  |
| 32H        | TAKEN_BRANCHES<br>(Counter 1)                              | Number of taken branches  |  |



**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name  | Description   | Comments   |
|------------|--|---|--|
| 33H        | D1_STARVATION_AND_FIFO_IS_EMPTY<br>(Counter 0)                                       | Number of times D1 stage cannot issue ANY instructions since the FIFO buffer is empty                           | The D1 stage can issue 0, 1, or 2 instructions per clock if those are available in an instructions FIFO buffer.  |
| 33H        | D1_STARVATION_AND_ONLY_ONE_INSTRUCTION_IN_FIFO<br>(Counter 1)                        | Number of times the D1 stage issues a single instruction (since the FIFO buffer had just one instruction ready) | The D1 stage can issue 0, 1, or 2 instructions per clock if those are available in an instructions FIFO buffer.<br>When combined with the previously defined events, Instruction Executed (16H) and Instruction Executed in the V-pipe (17H), this event enables the user to calculate the numbers of time pairing rules prevented issuing of two instructions.  |
| 34H        | MMX_INSTRUCTION_DATA_WRITES<br>(Counter 0)   | Number of data writes caused by MMX instructions  |  |
| 34H        | MMX_INSTRUCTION_DATA_WRITE_MISSES<br>(Counter 1)                                     | Number of data write misses caused by MMX instructions  |  |
| 35H        | PIPELINE_FLUSHES_DUE_TO_WRONG_BRANCH_PREDICTIONS<br>(Counter 0)                      | Number of pipeline flushes due to wrong branch predictions resolved in either the E-stage or the WB-stage       | The count includes any pipeline flush due to a branch that the pipeline did not follow correctly. It includes cases where a branch was not in the BTB, cases where a branch was in the BTB but was mispredicted, and cases where a branch was correctly predicted but to the wrong address.<br>Branches are resolved in either the Execute stage (E-stage) or the Writeback stage (WB-stage). In the later case, the misprediction penalty is larger by one clock. The difference between the 35H event count in counter 0 and counter 1 is the number of E-stage resolved branches. |
| 35H        | PIPELINE_FLUSHES_DUE_TO_WRONG_BRANCH_PREDICTIONS_RESOLVED_IN_WB-STAGE<br>(Counter 1) | Number of pipeline flushes due to wrong branch predictions resolved in the WB-stage                             | See note for event 35H (Counter 0).  |
| 36H        | MISALIGNED_DATA_MEMORY_REFERENCE_ON_MMX_INSTRUCTIONS<br>(Counter 0)                  | Number of misaligned data memory references when executing MMX instructions                                     |  |
| 36H        | PIPELINE_ISTALL_FOR_MMX_INSTRUCTION_DATA_MEMORY_READS<br>(Counter 1)                 | Number clocks during pipeline stalls caused by waits form MMX instruction data memory reads                     | T3:  |

**Table 19-41. Events That Can Be Counted with Pentium Processor Performance Monitoring Counters (Contd.)**

| Event Num. | Mnemonic Event Name  | Description   | Comments  |
|------------|--|---|---|
| 37H        | MISPREDICTED_OR_UNPREDICTED_RETURNS<br>(Counter 1)                               | Number of returns predicted incorrectly or not predicted at all   | The count is the difference between the total number of executed returns and the number of returns that were correctly predicted. Only RET instructions are counted (for example, IRET instructions are not counted).   |
| 37H        | PREDICTED_RETURNS<br>(Counter 1)   | Number of predicted returns (whether they are predicted correctly and incorrectly)  | Only RET instructions are counted (for example, IRET instructions are not counted).   |
| 38H        | MMX_MULTIPLY_UNIT_INTERLOCK<br>(Counter 0)                                       | Number of clocks the pipe is stalled since the destination of previous MMX multiply instruction is not ready yet  | The counter will not be incremented if there is another cause for a stall. For each occurrence of a multiply interlock, this event will be counted twice (if the stalled instruction comes on the next clock after the multiply) or by once (if the stalled instruction comes two clocks after the multiply). |
| 38H        | MOVD/MOVQ_STORE_STALL_DUE_TO_PREVIOUS_MMX_OPERATION<br>(Counter 1)               | Number of clocks a MOVD/MOVQ instruction store is stalled in D2 stage due to a previous MMX operation with a destination to be used in the store instruction. |   |
| 39H        | RETURNS<br>(Counter 0)   | Number of returns executed.   | Only RET instructions are counted; IRET instructions are not counted. Any exception taken on a RET instruction and any interrupt recognized by the processor on the instruction boundary prior to the execution of the RET instruction will also cause this counter to be incremented.                        |
| 39H        | Reserved   |   |   |
| 3AH        | BTB_FALSE_ENTRIES<br>(Counter 0)   | Number of false entries in the Branch Target Buffer   | False entries are causes for misprediction other than a wrong prediction.   |
| 3AH        | BTB_MISS_PREDICTION_ON_NOT-TAKEN_BRANCH<br>(Counter 1)                           | Number of times the BTB predicted a not-taken branch as taken   |   |
| 3BH        | FULL_WRITE_BUFFER_STALL_DURATION_WHILE_EXECUTING_MMX_INSTRUCTIONS<br>(Counter 0) | Number of clocks while the pipeline is stalled due to full write buffers while executing MMX instructions   |   |
| 3BH        | STALL_ON_MMX_INSTRUCTION_WRITE_TO_E_OR_M-STATE_LINE<br>(Counter 1)               | Number of clocks during stalls on MMX instructions writing to E- or M-state lines   |   |



### 13. Updates to Chapter 24, Volume 3B

Change bars show changes to Chapter 24 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2*.

-----

Changes to this chapter:

Update to Section 24.9.1 "Basic VM-Exit Information" to add XRSTORS and XSAVES to the list of exit qualifications.

### 24.1 OVERVIEW

A logical processor uses **virtual-machine control data structures (VMCSs)** while it is in VMX operation. These manage transitions into and out of VMX non-root operation (VM entries and VM exits) as well as processor behavior in VMX non-root operation. This structure is manipulated by the new instructions VMCLEAR, VMPTRLD, VMREAD, and VMWRITE.

A VMM can use a different VMCS for each virtual machine that it supports. For a virtual machine with multiple logical processors (virtual processors), the VMM can use a different VMCS for each virtual processor.

A logical processor associates a region in memory with each VMCS. This region is called the **VMCS region**.<sup>1</sup> Software references a specific VMCS using the 64-bit physical address of the region (a **VMCS pointer**). VMCS pointers must be aligned on a 4-KByte boundary (bits 11:0 must be zero). These pointers must not set bits beyond the processor's physical-address width.<sup>2,3</sup>

A logical processor may maintain a number of VMCSs that are active. The processor may optimize VMX operation by maintaining the state of an active VMCS in memory, on the processor, or both. At any given time, at most one of the active VMCSs is the **current VMCS**. (This document frequently uses the term "the VMCS" to refer to the current VMCS.) The VMLAUNCH, VMREAD, VMRESUME, and VMWRITE instructions operate only on the current VMCS.

The following items describe how a logical processor determines which VMCSs are active and which is current:

- The memory operand of the VMPTRLD instruction is the address of a VMCS. After execution of the instruction, that VMCS is both active and current on the logical processor. Any other VMCS that had been active remains so, but no other VMCS is current.
- The VMCS link pointer field in the current VMCS (see Section 24.4.2) is itself the address of a VMCS. If VM entry is performed successfully with the 1-setting of the "VMCS shadowing" VM-execution control, the VMCS referenced by the VMCS link pointer field becomes active on the logical processor. The identity of the current VMCS does not change.
- The memory operand of the VMCLEAR instruction is also the address of a VMCS. After execution of the instruction, that VMCS is neither active nor current on the logical processor. If the VMCS had been current on the logical processor, the logical processor no longer has a current VMCS.

The VMPTRST instruction stores the address of the logical processor's current VMCS into a specified memory location (it stores the value FFFFFFFF\_FFFFFFFFH if there is no current VMCS).

The **launch state** of a VMCS determines which VM-entry instruction should be used with that VMCS: the VMLAUNCH instruction requires a VMCS whose launch state is "clear"; the VMRESUME instruction requires a VMCS whose launch state is "launched". A logical processor maintains a VMCS's launch state in the corresponding VMCS region. The following items describe how a logical processor manages the launch state of a VMCS:

- If the launch state of the current VMCS is "clear", successful execution of the VMLAUNCH instruction changes the launch state to "launched".
- The memory operand of the VMCLEAR instruction is the address of a VMCS. After execution of the instruction, the launch state of that VMCS is "clear".
- There are no other ways to modify the launch state of a VMCS (it cannot be modified using VMWRITE) and there is no direct way to discover it (it cannot be read using VMREAD).

---

1. The amount of memory required for a VMCS region is at most 4 KBytes. The exact size is implementation specific and can be determined by consulting the VMX capability MSR IA32\_VMX\_BASIC to determine the size of the VMCS region (see Appendix A.1).

2. Software can determine a processor's physical-address width by executing CPUID with 80000008H in EAX. The physical-address width is returned in bits 7:0 of EAX.

3. If IA32\_VMX\_BASIC[48] is read as 1, these pointers must not set any bits in the range 63:32; see Appendix A.1.

Figure 24-1 illustrates the different states of a VMCS. It uses "X" to refer to the VMCS and "Y" to refer to any other VMCS. Thus: "VMPTRLD X" always makes X current and active; "VMPTRLD Y" always makes X not current (because it makes Y current); VMLAUNCH makes the launch state of X "launched" if X was current and its launch state was "clear"; and VMCLEAR X always makes X inactive and not current and makes its launch state "clear".

The figure does not illustrate operations that do not modify the VMCS state relative to these parameters (e.g., execution of VMPTRLD X when X is already current). Note that VMCLEAR X makes X "inactive, not current, and clear," even if X's current state is not defined (e.g., even if X has not yet been initialized). See Section 24.11.3.

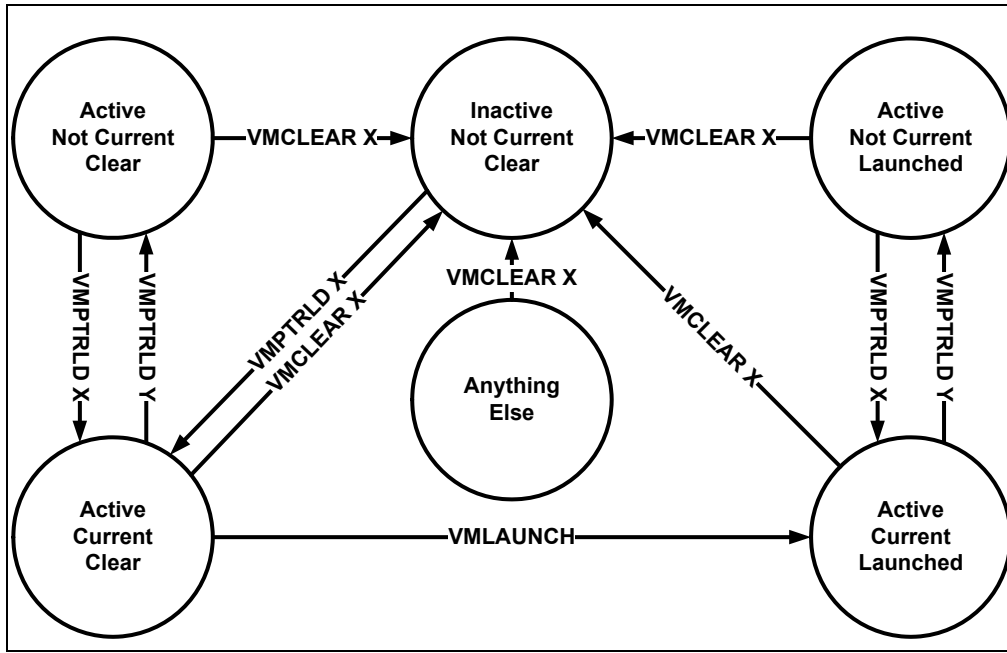


Figure 24-1. States of VMCS X

Because a shadow VMCS (see Section 24.10) cannot be used for VM entry, the launch state of a shadow VMCS is not meaningful. Figure 24-1 does not illustrate all the ways in which a shadow VMCS may be made active.

## 24.2 FORMAT OF THE VMCS REGION

A VMCS region comprises up to 4-KBytes.<sup>1</sup> The format of a VMCS region is given in Table 24-1.

Table 24-1. Format of the VMCS Region

| Byte Offset | Contents   |
|-------------|--|
| 0           | Bits 30:0: VMCS revision identifier<br>Bit 31: shadow-VMCS indicator (see Section 24.10) |
| 4           | VMX-abort indicator  |
| 8           | VMCS data (implementation-specific format)   |

The first 4 bytes of the VMCS region contain the **VMCS revision identifier** at bits 30:0.<sup>2</sup> Processors that maintain VMCS data in different formats (see below) use different VMCS revision identifiers. These identifiers enable soft-

1. The exact size is implementation specific and can be determined by consulting the VMX capability MSR IA32\_VMX\_BASIC to determine the size of the VMCS region (see Appendix A.1).

ware to avoid using a VMCS region formatted for one processor on a processor that uses a different format.<sup>1</sup> Bit 31 of this 4-byte region indicates whether the VMCS is a shadow VMCS (see Section 24.10).

Software should write the VMCS revision identifier to the VMCS region before using that region for a VMCS. The VMCS revision identifier is never written by the processor; VMPTRLD fails if its operand references a VMCS region whose VMCS revision identifier differs from that used by the processor. (VMPTRLD also fails if the shadow-VMCS indicator is 1 and the processor does not support the 1-setting of the “VMCS shadowing” VM-execution control; see Section 24.6.2.) Software can discover the VMCS revision identifier that a processor uses by reading the VMX capability MSR IA32\_VMX\_BASIC (see Appendix A.1).

Software should clear or set the shadow-VMCS indicator depending on whether the VMCS is to be an ordinary VMCS or a shadow VMCS (see Section 24.10). VMPTRLD fails if the shadow-VMCS indicator is set and the processor does not support the 1-setting of the “VMCS shadowing” VM-execution control. Software can discover support for this setting by reading the VMX capability MSR IA32\_VMX\_PROCBASED\_CTLS2 (see Appendix A.3.3).

The next 4 bytes of the VMCS region are used for the **VMX-abort indicator**. The contents of these bits do not control processor operation in any way. A logical processor writes a non-zero value into these bits if a VMX abort occurs (see Section 27.7). Software may also write into this field.

The remainder of the VMCS region is used for **VMCS data** (those parts of the VMCS that control VMX non-root operation and the VMX transitions). The format of these data is implementation-specific. VMCS data are discussed in Section 24.3 through Section 24.9. To ensure proper behavior in VMX operation, software should maintain the VMCS region and related structures (enumerated in Section 24.11.4) in writeback cacheable memory. Future implementations may allow or require a different memory type<sup>2</sup>. Software should consult the VMX capability MSR IA32\_VMX\_BASIC (see Appendix A.1).

## 24.3 ORGANIZATION OF VMCS DATA

The VMCS data are organized into six logical groups:

- **Guest-state area.** Processor state is saved into the guest-state area on VM exits and loaded from there on VM entries.
- **Host-state area.** Processor state is loaded from the host-state area on VM exits.
- **VM-execution control fields.** These fields control processor behavior in VMX non-root operation. They determine in part the causes of VM exits.
- **VM-exit control fields.** These fields control VM exits.
- **VM-entry control fields.** These fields control VM entries.
- **VM-exit information fields.** These fields receive information on VM exits and describe the cause and the nature of VM exits. On some processors, these fields are read-only.<sup>3</sup>

The VM-execution control fields, the VM-exit control fields, and the VM-entry control fields are sometimes referred to collectively as VMX controls.

---

2. Earlier versions of this manual specified that the VMCS revision identifier was a 32-bit field. For all processors produced prior to this change, bit 31 of the VMCS revision identifier was 0.

1. Logical processors that use the same VMCS revision identifier use the same size for VMCS regions.

2. Alternatively, software may map any of these regions or structures with the UC memory type. Doing so is strongly discouraged unless necessary as it will cause the performance of transitions using those structures to suffer significantly. In addition, the processor will continue to use the memory type reported in the VMX capability MSR IA32\_VMX\_BASIC with exceptions noted in Appendix A.1.

3. Software can discover whether these fields can be written by reading the VMX capability MSR IA32\_VMX\_MISC (see Appendix A.6).

## 24.4 GUEST-STATE AREA

This section describes fields contained in the guest-state area of the VMCS. As noted earlier, processor state is loaded from these fields on every VM entry (see Section 26.3.2) and stored into these fields on every VM exit (see Section 27.3).

### 24.4.1 Guest Register State

The following fields in the guest-state area correspond to processor registers:

- Control registers CR0, CR3, and CR4 (64 bits each; 32 bits on processors that do not support Intel 64 architecture).
- Debug register DR7 (64 bits; 32 bits on processors that do not support Intel 64 architecture).
- RSP, RIP, and RFLAGS (64 bits each; 32 bits on processors that do not support Intel 64 architecture).<sup>1</sup>
- The following fields for each of the registers CS, SS, DS, ES, FS, GS, LDTR, and TR:
  - Selector (16 bits).
  - Base address (64 bits; 32 bits on processors that do not support Intel 64 architecture). The base-address fields for CS, SS, DS, and ES have only 32 architecturally-defined bits; nevertheless, the corresponding VMCS fields have 64 bits on processors that support Intel 64 architecture.
  - Segment limit (32 bits). The limit field is always a measure in bytes.
  - Access rights (32 bits). The format of this field is given in Table 24-2 and detailed as follows:
    - The low 16 bits correspond to bits 23:8 of the upper 32 bits of a 64-bit segment descriptor. While bits 19:16 of code-segment and data-segment descriptors correspond to the upper 4 bits of the segment limit, the corresponding bits (bits 11:8) are reserved in this VMCS field.
    - Bit 16 indicates an unusable segment. Attempts to use such a segment fault except in 64-bit mode. In general, a segment register is unusable if it has been loaded with a null selector.<sup>2</sup>
    - Bits 31:17 are reserved.

**Table 24-2. Format of Access Rights**

| Bit Position(s) | Field  |
|-----------------|--|
| 3:0             | Segment type                                       |
| 4               | S — Descriptor type (0 = system; 1 = code or data) |
| 6:5             | DPL — Descriptor privilege level                   |
| 7               | P — Segment present                                |
| 11:8            | Reserved   |
| 12              | AVL — Available for use by system software         |

1. This chapter uses the notation RAX, RIP, RSP, RFLAGS, etc. for processor registers because most processors that support VMX operation also support Intel 64 architecture. For processors that do not support Intel 64 architecture, this notation refers to the 32-bit forms of those registers (EAX, EIP, ESP, EFLAGS, etc.). In a few places, notation such as EAX is used to refer specifically to lower 32 bits of the indicated register.
2. There are a few exceptions to this statement. For example, a segment with a non-null selector may be unusable following a task switch that fails after its commit point; see “Interrupt 10—Invalid TSS Exception (#TS)” in Section 6.14, “Exception and Interrupt Handling in 64-bit Mode,” of the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*. In contrast, the TR register is usable after processor reset despite having a null selector; see Table 10-1 in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.



Table 24-2. Format of Access Rights (Contd.)

| Bit Position(s) | Field   |
|-----------------|---|
| 13              | Reserved (except for CS)<br>L — 64-bit mode active (for CS only)      |
| 14              | D/B — Default operation size (0 = 16-bit segment; 1 = 32-bit segment) |
| 15              | G — Granularity   |
| 16              | Segment unusable (0 = usable; 1 = unusable)                           |
| 31:17           | Reserved  |

The base address, segment limit, and access rights compose the “hidden” part (or “descriptor cache”) of each segment register. These data are included in the VMCS because it is possible for a segment register’s descriptor cache to be inconsistent with the segment descriptor in memory (in the GDT or the LDT) referenced by the segment register’s selector.

The value of the DPL field for SS is always equal to the logical processor’s current privilege level (CPL).<sup>1</sup>

- The following fields for each of the registers GDTR and IDTR:
  - Base address (64 bits; 32 bits on processors that do not support Intel 64 architecture).
  - Limit (32 bits). The limit fields contain 32 bits even though these fields are specified as only 16 bits in the architecture.
- The following MSRs:
  - IA32\_DEBUGCTL (64 bits)
  - IA32\_SYSENTER\_CS (32 bits)
  - IA32\_SYSENTER\_ESP and IA32\_SYSENTER\_EIP (64 bits; 32 bits on processors that do not support Intel 64 architecture)
  - IA32\_PERF\_GLOBAL\_CTRL (64 bits). This field is supported only on processors that support the 1-setting of the “load IA32\_PERF\_GLOBAL\_CTRL” VM-entry control.
  - IA32\_PAT (64 bits). This field is supported only on processors that support either the 1-setting of the “load IA32\_PAT” VM-entry control or that of the “save IA32\_PAT” VM-exit control.
  - IA32\_EFER (64 bits). This field is supported only on processors that support either the 1-setting of the “load IA32\_EFER” VM-entry control or that of the “save IA32\_EFER” VM-exit control.
  - IA32\_BNDCFGS (64 bits). This field is supported only on processors that support either the 1-setting of the “load IA32\_BNDCFGS” VM-entry control or that of the “clear IA32\_BNDCFGS” VM-exit control.
- The register SMBASE (32 bits). This register contains the base address of the logical processor’s SMRAM image.

## 24.4.2 Guest Non-Register State

In addition to the register state described in Section 24.4.1, the guest-state area includes the following fields that characterize guest state but which do not correspond to processor registers:

- **Activity state** (32 bits). This field identifies the logical processor’s activity state. When a logical processor is executing instructions normally, it is in the **active state**. Execution of certain instructions and the occurrence of certain events may cause a logical processor to transition to an **inactive state** in which it ceases to execute instructions.

The following activity states are defined:<sup>2</sup>

- **0: Active**. The logical processor is executing instructions normally.

1. In protected mode, CPL is also associated with the RPL field in the CS selector. However, the RPL fields are not meaningful in real-address mode or in virtual-8086 mode.

- 1: **HLT**. The logical processor is inactive because it executed the HLT instruction.
- 2: **Shutdown**. The logical processor is inactive because it incurred a **triple fault**<sup>1</sup> or some other serious error.
- 3: **Wait-for-SIPI**. The logical processor is inactive because it is waiting for a startup-IPI (SIPI).

Future processors may include support for other activity states. Software should read the VMX capability MSR IA32\_VMX\_MISC (see Appendix A.6) to determine what activity states are supported.

- **Interruptibility state** (32 bits). The IA-32 architecture includes features that permit certain events to be blocked for a period of time. This field contains information about such blocking. Details and the format of this field are given in Table 24-3.

**Table 24-3. Format of Interruptibility State**

| Bit Position(s) | Bit Name             | Notes   |
|-----------------|----------------------|---|
| 0               | Blocking by STI      | See the “STI—Set Interrupt Flag” section in Chapter 4 of the <i>Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B</i> .<br>Execution of STI with RFLAGS.IF = 0 blocks interrupts (and, optionally, other events) for one instruction after its execution. Setting this bit indicates that this blocking is in effect.  |
| 1               | Blocking by MOV SS   | See the “MOV—Move a Value from the Stack” from Chapter 4 of the <i>Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B</i> , and “POP—Pop a Value from the Stack” from Chapter 4 of the <i>Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B</i> , and Section 6.8.3 in the <i>Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A</i> .<br>Execution of a MOV to SS or a POP to SS blocks interrupts for one instruction after its execution. In addition, certain debug exceptions are inhibited between a MOV to SS or a POP to SS and a subsequent instruction. Setting this bit indicates that the blocking of all these events is in effect. This document uses the term “blocking by MOV SS,” but it applies equally to POP SS. |
| 2               | Blocking by SMI      | See Section 34.2. System-management interrupts (SMIs) are disabled while the processor is in system-management mode (SMM). Setting this bit indicates that blocking of SMIs is in effect.   |
| 3               | Blocking by NMI      | See Section 6.7.1 in the <i>Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A</i> and Section 34.8.<br>Delivery of a non-maskable interrupt (NMI) or a system-management interrupt (SMI) blocks subsequent NMIs until the next execution of IRET. See Section 25.3 for how this behavior of IRET may change in VMX non-root operation. Setting this bit indicates that blocking of NMIs is in effect. Clearing this bit does not imply that NMIs are not (temporarily) blocked for other reasons.<br>If the “virtual NMIs” VM-execution control (see Section 24.6.1) is 1, this bit does not control the blocking of NMIs. Instead, it refers to “virtual-NMI blocking” (the fact that guest software is not ready for an NMI).  |
| 4               | Enclave interruption | A VM exit saves this bit as 1 to indicate that the VM exit was incident to enclave mode.  |
| 31:5            | Reserved             | VM entry will fail if these bits are not 0. See Section 26.3.1.5.   |

- **Pending debug exceptions** (64 bits; 32 bits on processors that do not support Intel 64 architecture). IA-32 processors may recognize one or more debug exceptions without immediately delivering them.<sup>2</sup> This field contains information about such exceptions. This field is described in Table 24-4.

2. Execution of the MWAIT instruction may put a logical processor into an inactive state. However, this VMCS field never reflects this state. See Section 27.1.

1. A triple fault occurs when a logical processor encounters an exception while attempting to deliver a double fault.

**Table 24-4. Format of Pending-Debug-Exceptions**

| Bit Position(s) | Bit Name           | Notes  |
|-----------------|--------------------|--|
| 3:0             | B3 - B0            | When set, each of these bits indicates that the corresponding breakpoint condition was met. Any of these bits may be set even if the corresponding enabling bit in DR7 is not set.   |
| 11:4            | Reserved           | VM entry fails if these bits are not 0. See Section 26.3.1.5.  |
| 12              | Enabled breakpoint | When set, this bit indicates that at least one data or I/O breakpoint was met and was enabled in DR7.  |
| 13              | Reserved           | VM entry fails if this bit is not 0. See Section 26.3.1.5.   |
| 14              | BS                 | When set, this bit indicates that a debug exception would have been triggered by single-step execution mode.   |
| 15              | Reserved           | VM entry fails if this bit is not 0. See Section 26.3.1.5.   |
| 16              | RTM                | When set, this bit indicates that a debug exception (#DB) or a breakpoint exception (#BP) occurred inside an RTM region while advanced debugging of RTM transactional regions was enabled (see Section 16.3.7, "RTM-Enabled Debugger Support," of <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1</i> ). <sup>1</sup> |
| 63:17           | Reserved           | VM entry fails if these bits are not 0. See Section 26.3.1.5. Bits 63:32 exist only on processors that support Intel 64 architecture.  |

**NOTES:**

1. In general, the format of this field matches that of DR6. However, DR6 **clears** bit 16 to indicate an RTM-related exception, while this field **sets** the bit to indicate that condition.

- **VMCS link pointer** (64 bits). If the "VMCS shadowing" VM-execution control is 1, the VMREAD and VMWRITE instructions access the VMCS referenced by this pointer (see Section 24.10). Otherwise, software should set this field to FFFFFFFF\_FFFFFFFFH to avoid VM-entry failures (see Section 26.3.1.5).
- **VMX-preemption timer value** (32 bits). This field is supported only on processors that support the 1-setting of the "activate VMX-preemption timer" VM-execution control. This field contains the value that the VMX-preemption timer will use following the next VM entry with that setting. See Section 25.5.1 and Section 26.6.4.
- **Page-directory-pointer-table entries** (PDPTes; 64 bits each). These four (4) fields (PDPTE0, PDPTE1, PDPTE2, and PDPTE3) are supported only on processors that support the 1-setting of the "enable EPT" VM-execution control. They correspond to the PDPTes referenced by CR3 when PAE paging is in use (see Section 4.4 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*). They are used only if the "enable EPT" VM-execution control is 1.
- **Guest interrupt status** (16 bits). This field is supported only on processors that support the 1-setting of the "virtual-interrupt delivery" VM-execution control. It characterizes part of the guest's virtual-APIC state and does not correspond to any processor or APIC registers. It comprises two 8-bit subfields:
  - **Requesting virtual interrupt (RVI)**. This is the low byte of the guest interrupt status. The processor treats this value as the vector of the highest priority virtual interrupt that is requesting service. (The value 0 implies that there is no such interrupt.)
  - **Servicing virtual interrupt (SVI)**. This is the high byte of the guest interrupt status. The processor treats this value as the vector of the highest priority virtual interrupt that is in service. (The value 0 implies that there is no such interrupt.)

2. For example, execution of a MOV to SS or a POP to SS may inhibit some debug exceptions for one instruction. See Section 6.8.3 of *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*. In addition, certain events incident to an instruction (for example, an INIT signal) may take priority over debug traps generated by that instruction. See Table 6-2 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.

See Chapter 29 for more information on the use of this field.

- **PML index** (16 bits). This field is supported only on processors that support the 1-setting of the “enable PML” VM-execution control. It contains the logical index of the next entry in the page-modification log. Because the page-modification log comprises 512 entries, the PML index is typically a value in the range 0–511. Details of the page-modification log and use of the PML index are given in Section 28.2.5.

## 24.5 HOST-STATE AREA

This section describes fields contained in the host-state area of the VMCS. As noted earlier, processor state is loaded from these fields on every VM exit (see Section 27.5).

All fields in the host-state area correspond to processor registers:

- CR0, CR3, and CR4 (64 bits each; 32 bits on processors that do not support Intel 64 architecture).
- RSP and RIP (64 bits each; 32 bits on processors that do not support Intel 64 architecture).
- Selector fields (16 bits each) for the segment registers CS, SS, DS, ES, FS, GS, and TR. There is no field in the host-state area for the LDTR selector.
- Base-address fields for FS, GS, TR, GDTR, and IDTR (64 bits each; 32 bits on processors that do not support Intel 64 architecture).
- The following MSRs:
  - IA32\_SYSENTER\_CS (32 bits)
  - IA32\_SYSENTER\_ESP and IA32\_SYSENTER\_EIP (64 bits; 32 bits on processors that do not support Intel 64 architecture).
  - IA32\_PERF\_GLOBAL\_CTRL (64 bits). This field is supported only on processors that support the 1-setting of the “load IA32\_PERF\_GLOBAL\_CTRL” VM-exit control.
  - IA32\_PAT (64 bits). This field is supported only on processors that support the 1-setting of the “load IA32\_PAT” VM-exit control.
  - IA32\_EFER (64 bits). This field is supported only on processors that support the 1-setting of the “load IA32\_EFER” VM-exit control.

In addition to the state identified here, some processor state components are loaded with fixed values on every VM exit; there are no fields corresponding to these components in the host-state area. See Section 27.5 for details of how state is loaded on VM exits.

## 24.6 VM-EXECUTION CONTROL FIELDS

The VM-execution control fields govern VMX non-root operation. These are described in Section 24.6.1 through Section 24.6.8.

### 24.6.1 Pin-Based VM-Execution Controls

The pin-based VM-execution controls constitute a 32-bit vector that governs the handling of asynchronous events (for example: interrupts).<sup>1</sup> Table 24-5 lists the controls. See Chapter 27 for how these controls affect processor behavior in VMX non-root operation.

---

1. Some asynchronous events cause VM exits regardless of the settings of the pin-based VM-execution controls (see Section 25.2).

**Table 24-5. Definitions of Pin-Based VM-Execution Controls**

| Bit Position(s) | Name                          | Description  |
|-----------------|-------------------------------|--|
| 0               | External-interrupt exiting    | If this control is 1, external interrupts cause VM exits. Otherwise, they are delivered normally through the guest interrupt-descriptor table (IDT). If this control is 1, the value of RFLAGS.IF does not affect interrupt blocking.  |
| 3               | NMI exiting                   | If this control is 1, non-maskable interrupts (NMIs) cause VM exits. Otherwise, they are delivered normally using descriptor 2 of the IDT. This control also determines interactions between IRET and blocking by NMI (see Section 25.3).                                      |
| 5               | Virtual NMIs                  | If this control is 1, NMIs are never blocked and the “blocking by NMI” bit (bit 3) in the interruptibility-state field indicates “virtual-NMI blocking” (see Table 24-3). This control also interacts with the “NMI-window exiting” VM-execution control (see Section 24.6.2). |
| 6               | Activate VMX-preemption timer | If this control is 1, the VMX-preemption timer counts down in VMX non-root operation; see Section 25.5.1. A VM exit occurs when the timer counts down to zero; see Section 25.2.   |
| 7               | Process posted interrupts     | If this control is 1, the processor treats interrupts with the posted-interrupt notification vector (see Section 24.6.8) specially, updating the virtual-APIC page with posted-interrupt requests (see Section 29.6).  |

All other bits in this field are reserved, some to 0 and some to 1. Software should consult the VMX capability MSRs IA32\_VMX\_PINBASED\_CTLs and IA32\_VMX\_TRUE\_PINBASED\_CTLs (see Appendix A.3.1) to determine how to set reserved bits. Failure to set reserved bits properly causes subsequent VM entries to fail (see Section 26.2.1.1).

The first processors to support the virtual-machine extensions supported only the 1-settings of bits 1, 2, and 4. The VMX capability MSR IA32\_VMX\_PINBASED\_CTLs will always report that these bits must be 1. Logical processors that support the 0-settings of any of these bits will support the VMX capability MSR IA32\_VMX\_TRUE\_PINBASED\_CTLs MSR, and software should consult this MSR to discover support for the 0-settings of these bits. Software that is not aware of the functionality of any one of these bits should set that bit to 1.

## 24.6.2 Processor-Based VM-Execution Controls

The processor-based VM-execution controls constitute two 32-bit vectors that govern the handling of synchronous events, mainly those caused by the execution of specific instructions.<sup>1</sup> These are the **primary processor-based VM-execution controls** and the **secondary processor-based VM-execution controls**.

Table 24-6 lists the primary processor-based VM-execution controls. See Chapter 25 for more details of how these controls affect processor behavior in VMX non-root operation.

**Table 24-6. Definitions of Primary Processor-Based VM-Execution Controls**

| Bit Position(s) | Name                     | Description  |
|-----------------|--------------------------|--|
| 2               | Interrupt-window exiting | If this control is 1, a VM exit occurs at the beginning of any instruction if RFLAGS.IF = 1 and there are no other blocking of interrupts (see Section 24.4.2).  |
| 3               | Use TSC offsetting       | This control determines whether executions of RDTSC, executions of RDTSCP, and executions of RDMSR that read from the IA32_TIME_STAMP_COUNTER MSR return a value modified by the TSC offset field (see Section 24.6.5 and Section 25.3). |
| 7               | HLT exiting              | This control determines whether executions of HLT cause VM exits.  |
| 9               | INVLPG exiting           | This determines whether executions of INVLPG cause VM exits.   |
| 10              | MWAIT exiting            | This control determines whether executions of MWAIT cause VM exits.  |
| 11              | RDPIC exiting            | This control determines whether executions of RDPIC cause VM exits.  |
| 12              | RDTSC exiting            | This control determines whether executions of RDTSC and RDTSCP cause VM exits.   |

1. Some instructions cause VM exits regardless of the settings of the processor-based VM-execution controls (see Section 25.1.2), as do task switches (see Section 25.2).

**Table 24-6. Definitions of Primary Processor-Based VM-Execution Controls (Contd.)**

| Bit Position(s) | Name                        | Description   |
|-----------------|-----------------------------|---|
| 15              | CR3-load exiting            | In conjunction with the CR3-target controls (see Section 24.6.7), this control determines whether executions of MOV to CR3 cause VM exits. See Section 25.1.3.<br>The first processors to support the virtual-machine extensions supported only the 1-setting of this control.  |
| 16              | CR3-store exiting           | This control determines whether executions of MOV from CR3 cause VM exits.<br>The first processors to support the virtual-machine extensions supported only the 1-setting of this control.  |
| 19              | CR8-load exiting            | This control determines whether executions of MOV to CR8 cause VM exits.  |
| 20              | CR8-store exiting           | This control determines whether executions of MOV from CR8 cause VM exits.  |
| 21              | Use TPR shadow              | Setting this control to 1 enables TPR virtualization and other APIC-virtualization features. See Chapter 29.  |
| 22              | NMI-window exiting          | If this control is 1, a VM exit occurs at the beginning of any instruction if there is no virtual-NMI blocking (see Section 24.4.2).  |
| 23              | MOV-DR exiting              | This control determines whether executions of MOV DR cause VM exits.  |
| 24              | Unconditional I/O exiting   | This control determines whether executions of I/O instructions (IN, INS/INSB/INSW/INSD, OUT, and OUTS/OUTSB/OUTSW/OUTSD) cause VM exits.  |
| 25              | Use I/O bitmaps             | This control determines whether I/O bitmaps are used to restrict executions of I/O instructions (see Section 24.6.4 and Section 25.1.3).<br>For this control, “0” means “do not use I/O bitmaps” and “1” means “use I/O bitmaps.” If the I/O bitmaps are used, the setting of the “unconditional I/O exiting” control is ignored.                   |
| 27              | Monitor trap flag           | If this control is 1, the monitor trap flag debugging feature is enabled. See Section 25.5.2.   |
| 28              | Use MSR bitmaps             | This control determines whether MSR bitmaps are used to control execution of the RDMSR and WRMSR instructions (see Section 24.6.9 and Section 25.1.3).<br>For this control, “0” means “do not use MSR bitmaps” and “1” means “use MSR bitmaps.” If the MSR bitmaps are not used, all executions of the RDMSR and WRMSR instructions cause VM exits. |
| 29              | MONITOR exiting             | This control determines whether executions of MONITOR cause VM exits.   |
| 30              | PAUSE exiting               | This control determines whether executions of PAUSE cause VM exits.   |
| 31              | Activate secondary controls | This control determines whether the secondary processor-based VM-execution controls are used. If this control is 0, the logical processor operates as if all the secondary processor-based VM-execution controls were also 0.   |

All other bits in this field are reserved, some to 0 and some to 1. Software should consult the VMX capability MSRs IA32\_VMX\_PROCBASED\_CTLs and IA32\_VMX\_TRUE\_PROCBASED\_CTLs (see Appendix A.3.2) to determine how to set reserved bits. Failure to set reserved bits properly causes subsequent VM entries to fail (see Section 26.2.1.1).

The first processors to support the virtual-machine extensions supported only the 1-settings of bits 1, 4–6, 8, 13–16, and 26. The VMX capability MSR IA32\_VMX\_PROCBASED\_CTLs will always report that these bits must be 1. Logical processors that support the 0-settings of any of these bits will support the VMX capability MSR IA32\_VMX\_TRUE\_PROCBASED\_CTLs MSR, and software should consult this MSR to discover support for the 0-settings of these bits. Software that is not aware of the functionality of any one of these bits should set that bit to 1.

Bit 31 of the primary processor-based VM-execution controls determines whether the secondary processor-based VM-execution controls are used. If that bit is 0, VM entry and VMX non-root operation function as if all the secondary processor-based VM-execution controls were 0. Processors that support only the 0-setting of bit 31 of the primary processor-based VM-execution controls do not support the secondary processor-based VM-execution controls.

Table 24-7 lists the secondary processor-based VM-execution controls. See Chapter 25 for more details of how these controls affect processor behavior in VMX non-root operation.

**Table 24-7. Definitions of Secondary Processor-Based VM-Execution Controls**

| Bit Position(s) | Name                               | Description  |
|-----------------|------------------------------------|--|
| 0               | Virtualize APIC accesses           | If this control is 1, the logical processor treats specially accesses to the page with the APIC-access address. See Section 29.4.  |
| 1               | Enable EPT                         | If this control is 1, extended page tables (EPT) are enabled. See Section 28.2.  |
| 2               | Descriptor-table exiting           | This control determines whether executions of LGDT, LIDT, LLDT, LTR, SGDT, SIDT, SLDT, and STR cause VM exits.   |
| 3               | Enable RDTSCP                      | If this control is 0, any execution of RDTSCP causes an invalid-opcode exception (#UD).  |
| 4               | Virtualize x2APIC mode             | If this control is 1, the logical processor treats specially RDMSR and WRMSR to APIC MSRs (in the range 800H–8FFH). See Section 29.5.  |
| 5               | Enable VPID                        | If this control is 1, cached translations of linear addresses are associated with a virtual-processor identifier (VPID). See Section 28.1.   |
| 6               | WBINVD exiting                     | This control determines whether executions of WBINVD cause VM exits.   |
| 7               | Unrestricted guest                 | This control determines whether guest software may run in unpagged protected mode or in real-address mode.   |
| 8               | APIC-register virtualization       | If this control is 1, the logical processor virtualizes certain APIC accesses. See Section 29.4 and Section 29.5.  |
| 9               | Virtual-interrupt delivery         | This control enables the evaluation and delivery of pending virtual interrupts as well as the emulation of writes to the APIC registers that control interrupt prioritization.   |
| 10              | PAUSE-loop exiting                 | This control determines whether a series of executions of PAUSE can cause a VM exit (see Section 24.6.13 and Section 25.1.3).  |
| 11              | RDRAND exiting                     | This control determines whether executions of RDRAND cause VM exits.   |
| 12              | Enable INVPCID                     | If this control is 0, any execution of INVPCID causes a #UD.   |
| 13              | Enable VM functions                | Setting this control to 1 enables use of the VMFUNC instruction in VMX non-root operation. See Section 25.5.5.   |
| 14              | VMCS shadowing                     | If this control is 1, executions of VMREAD and VMWRITE in VMX non-root operation may access a shadow VMCS (instead of causing VM exits). See Section 24.10 and Section 30.3.   |
| 15              | Enable ENCLS exiting               | If this control is 1, executions of ENCLS consult the ENCLS-exiting bitmap to determine whether the instruction causes a VM exit. See Section 24.6.16 and Section 25.1.3.  |
| 16              | RDSEED exiting                     | This control determines whether executions of RDSEED cause VM exits.   |
| 17              | Enable PML                         | If this control is 1, an access to a guest-physical address that sets an EPT dirty bit first adds an entry to the page-modification log. See Section 28.2.5.   |
| 18              | EPT-violation #VE                  | If this control is 1, EPT violations may cause virtualization exceptions (#VE) instead of VM exits. See Section 25.5.6.  |
| 19              | Conceal VMX from PT                | If this control is 1, Intel Processor Trace suppresses from PIPs an indication that the processor was in VMX non-root operation and omits a VMCS packet from any PSB+ produced in VMX non-root operation (see Chapter 35).                   |
| 20              | Enable XSAVES/XRSTORS              | If this control is 0, any execution of XSAVES or XRSTORS causes a #UD.   |
| 22              | Mode-based execute control for EPT | If this control is 1, EPT execute permissions are based on whether the linear address being accessed is supervisor mode or user mode. See Chapter 28.  |
| 25              | Use TSC scaling                    | This control determines whether executions of RDTSC, executions of RDTSCP, and executions of RDMSR that read from the IA32_TIME_STAMP_COUNTER MSR return a value modified by the TSC multiplier field (see Section 24.6.5 and Section 25.3). |

All other bits in this field are reserved to 0. Software should consult the VMX capability MSR IA32\_VMX\_PROCBASED\_CTL2 (see Appendix A.3.3) to determine which bits may be set to 1. Failure to clear reserved bits causes subsequent VM entries to fail (see Section 26.2.1.1).

### 24.6.3 Exception Bitmap

The **exception bitmap** is a 32-bit field that contains one bit for each exception. When an exception occurs, its vector is used to select a bit in this field. If the bit is 1, the exception causes a VM exit. If the bit is 0, the exception is delivered normally through the IDT, using the descriptor corresponding to the exception's vector.

Whether a page fault (exception with vector 14) causes a VM exit is determined by bit 14 in the exception bitmap as well as the error code produced by the page fault and two 32-bit fields in the VMCS (the **page-fault error-code mask** and **page-fault error-code match**). See Section 25.2 for details.

### 24.6.4 I/O-Bitmap Addresses

The VM-execution control fields include the 64-bit physical addresses of **I/O bitmaps A** and **B** (each of which are 4 KBytes in size). I/O bitmap A contains one bit for each I/O port in the range 0000H through 7FFFH; I/O bitmap B contains bits for ports in the range 8000H through FFFFH.

A logical processor uses these bitmaps if and only if the "use I/O bitmaps" control is 1. If the bitmaps are used, execution of an I/O instruction causes a VM exit if any bit in the I/O bitmaps corresponding to a port it accesses is 1. See Section 25.1.3 for details. If the bitmaps are used, their addresses must be 4-KByte aligned.

### 24.6.5 Time-Stamp Counter Offset and Multiplier

The VM-execution control fields include a 64-bit **TSC-offset** field. If the "RDTSC exiting" control is 0 and the "use TSC offsetting" control is 1, this field controls executions of the RDTSC and RDTSCP instructions. It also controls executions of the RDMSR instruction that read from the IA32\_TIME\_STAMP\_COUNTER MSR. For all of these, the value of the TSC offset is added to the value of the time-stamp counter, and the sum is returned to guest software in EDX:EAX.

Processors that support the 1-setting of the "use TSC scaling" control also support a 64-bit **TSC-multiplier** field. If this control is 1 (and the "RDTSC exiting" control is 0 and the "use TSC offsetting" control is 1), this field also affects the executions of the RDTSC, RDTSCP, and RDMSR instructions identified above. Specifically, the contents of the time-stamp counter is first multiplied by the TSC multiplier before adding the TSC offset.

See Chapter 27 for a detailed treatment of the behavior of RDTSC, RDTSCP, and RDMSR in VMX non-root operation.

### 24.6.6 Guest/Host Masks and Read Shadows for CR0 and CR4

VM-execution control fields include **guest/host masks** and **read shadows** for the CR0 and CR4 registers. These fields control executions of instructions that access those registers (including CLTS, LMSW, MOV CR, and SMSW). They are 64 bits on processors that support Intel 64 architecture and 32 bits on processors that do not.

In general, bits set to 1 in a guest/host mask correspond to bits "owned" by the host:

- Guest attempts to set them (using CLTS, LMSW, or MOV to CR) to values differing from the corresponding bits in the corresponding read shadow cause VM exits.
- Guest reads (using MOV from CR or SMSW) return values for these bits from the corresponding read shadow.

Bits cleared to 0 correspond to bits "owned" by the guest; guest attempts to modify them succeed and guest reads return values for these bits from the control register itself.

See Chapter 27 for details regarding how these fields affect VMX non-root operation.



## 24.6.7 CR3-Target Controls

The VM-execution control fields include a set of 4 **CR3-target values** and a **CR3-target count**. The CR3-target values each have 64 bits on processors that support Intel 64 architecture and 32 bits on processors that do not. The CR3-target count has 32 bits on all processors.

An execution of MOV to CR3 in VMX non-root operation does not cause a VM exit if its source operand matches one of these values. If the CR3-target count is  $n$ , only the first  $n$  CR3-target values are considered; if the CR3-target count is 0, MOV to CR3 always causes a VM exit

There are no limitations on the values that can be written for the CR3-target values. VM entry fails (see Section 26.2) if the CR3-target count is greater than 4.

Future processors may support a different number of CR3-target values. Software should read the VMX capability MSR IA32\_VMX\_MISC (see Appendix A.6) to determine the number of values supported.

## 24.6.8 Controls for APIC Virtualization

There are three mechanisms by which software accesses registers of the logical processor's local APIC:

- If the local APIC is in xAPIC mode, it can perform memory-mapped accesses to addresses in the 4-KByte page referenced by the physical address in the IA32\_APIC\_BASE MSR (see Section 10.4.4, "Local APIC Status and Location" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A* and *Intel® 64 Architecture Processor Topology Enumeration*).<sup>1</sup>
- If the local APIC is in x2APIC mode, it can access the local APIC's registers using the RDMSR and WRMSR instructions (see *Intel® 64 Architecture Processor Topology Enumeration*).
- In 64-bit mode, it can access the local APIC's task-priority register (TPR) using the MOV CR8 instruction.

There are five processor-based VM-execution controls (see Section 24.6.2) that control such accesses. There are "use TPR shadow", "virtualize APIC accesses", "virtualize x2APIC mode", "virtual-interrupt delivery", and "APIC-register virtualization". These controls interact with the following fields:

- **APIC-access address** (64 bits). This field contains the physical address of the 4-KByte **APIC-access page**. If the "virtualize APIC accesses" VM-execution control is 1, access to this page may cause VM exits or be virtualized by the processor. See Section 29.4.

The APIC-access address exists only on processors that support the 1-setting of the "virtualize APIC accesses" VM-execution control.

- **Virtual-APIC address** (64 bits). This field contains the physical address of the 4-KByte **virtual-APIC page**. The processor uses the virtual-APIC page to virtualize certain accesses to APIC registers and to manage virtual interrupts; see Chapter 29.

Depending on the setting of the controls indicated earlier, the virtual-APIC page may be accessed by the following operations:

- The MOV CR8 instructions (see Section 29.3).
- Accesses to the APIC-access page if, in addition, the "virtualize APIC accesses" VM-execution control is 1 (see Section 29.4).
- The RDMSR and WRMSR instructions if, in addition, the value of ECX is in the range 800H–8FFH (indicating an APIC MSR) and the "virtualize x2APIC mode" VM-execution control is 1 (see Section 29.5).

If the "use TPR shadow" VM-execution control is 1, VM entry ensures that the virtual-APIC address is 4-KByte aligned. The virtual-APIC address exists only on processors that support the 1-setting of the "use TPR shadow" VM-execution control.

- **TPR threshold** (32 bits). Bits 3:0 of this field determine the threshold below which bits 7:4 of VTPR (see Section 29.1.1) cannot fall. If the "virtual-interrupt delivery" VM-execution control is 0, a VM exit occurs after an operation (e.g., an execution of MOV to CR8) that reduces the value of those bits below the TPR threshold. See Section 29.1.2.

---

1. If the local APIC does not support x2APIC mode, it is always in xAPIC mode.

The TPR threshold exists only on processors that support the 1-setting of the “use TPR shadow” VM-execution control.

- **EOI -exit bitmap** (4 fields; 64 bits each). These fields are supported only on processors that support the 1-setting of the “virtual-interrupt delivery” VM-execution control. They are used to determine which virtualized writes to the APIC’s EOI register cause VM exits:
  - EOI\_EXIT0 contains bits for vectors from 0 (bit 0) to 63 (bit 63).
  - EOI\_EXIT1 contains bits for vectors from 64 (bit 0) to 127 (bit 63).
  - EOI\_EXIT2 contains bits for vectors from 128 (bit 0) to 191 (bit 63).
  - EOI\_EXIT3 contains bits for vectors from 192 (bit 0) to 255 (bit 63).
 See Section 29.1.4 for more information on the use of this field.
- **Posted-interrupt notification vector** (16 bits). This field is supported only on processors that support the 1-setting of the “process posted interrupts” VM-execution control. Its low 8 bits contain the interrupt vector that is used to notify a logical processor that virtual interrupts have been posted. See Section 29.6 for more information on the use of this field.
- **Posted-interrupt descriptor address** (64 bits). This field is supported only on processors that support the 1-setting of the “process posted interrupts” VM-execution control. It is the physical address of a 64-byte aligned posted interrupt descriptor. See Section 29.6 for more information on the use of this field.

## 24.6.9 MSR-Bitmap Address

On processors that support the 1-setting of the “use MSR bitmaps” VM-execution control, the VM-execution control fields include the 64-bit physical address of four contiguous **MSR bitmaps**, which are each 1-KByte in size. This field does not exist on processors that do not support the 1-setting of that control. The four bitmaps are:

- **Read bitmap for low MSRs** (located at the MSR-bitmap address). This contains one bit for each MSR address in the range 00000000H to 00001FFFH. The bit determines whether an execution of RDMSR applied to that MSR causes a VM exit.
- **Read bitmap for high MSRs** (located at the MSR-bitmap address plus 1024). This contains one bit for each MSR address in the range C0000000H to C0001FFFH. The bit determines whether an execution of RDMSR applied to that MSR causes a VM exit.
- **Write bitmap for low MSRs** (located at the MSR-bitmap address plus 2048). This contains one bit for each MSR address in the range 00000000H to 00001FFFH. The bit determines whether an execution of WRMSR applied to that MSR causes a VM exit.
- **Write bitmap for high MSRs** (located at the MSR-bitmap address plus 3072). This contains one bit for each MSR address in the range C0000000H to C0001FFFH. The bit determines whether an execution of WRMSR applied to that MSR causes a VM exit.

A logical processor uses these bitmaps if and only if the “use MSR bitmaps” control is 1. If the bitmaps are used, an execution of RDMSR or WRMSR causes a VM exit if the value of RCX is in neither of the ranges covered by the bitmaps or if the appropriate bit in the MSR bitmaps (corresponding to the instruction and the RCX value) is 1. See Section 25.1.3 for details. If the bitmaps are used, their address must be 4-KByte aligned.

## 24.6.10 Executive-VMCS Pointer

The executive-VMCS pointer is a 64-bit field used in the dual-monitor treatment of system-management interrupts (SMIs) and system-management mode (SMM). SMM VM exits save this field as described in Section 34.15.2. VM entries that return from SMM use this field as described in Section 34.15.4.

### 24.6.11 Extended-Page-Table Pointer (EPTP)

The **extended-page-table pointer** (EPTP) contains the address of the base of EPT PML4 table (see Section 28.2.2), as well as other EPT configuration information. The format of this field is shown in Table 24-8.

**Table 24-8. Format of Extended-Page-Table Pointer**

| Bit Position(s) | Field  |
|-----------------|--|
| 2:0             | EPT paging-structure memory type (see Section 28.2.6):<br>0 = Uncacheable (UC)<br>6 = Write-back (WB)<br>Other values are reserved. <sup>1</sup> |
| 5:3             | This value is 1 less than the EPT page-walk length (see Section 28.2.2)  |
| 6               | Setting this control to 1 enables accessed and dirty flags for EPT (see Section 28.2.4) <sup>2</sup>   |
| 11:7            | Reserved   |
| N-1:12          | Bits N-1:12 of the physical address of the 4-KByte aligned EPT PML4 table <sup>3</sup>   |
| 63:N            | Reserved   |

#### NOTES:

1. Software should read the VMX capability MSR IA32\_VMX\_EPT\_VPID\_CAP (see Appendix A.10) to determine what EPT paging-structure memory types are supported.
2. Not all processors support accessed and dirty flags for EPT. Software should read the VMX capability MSR IA32\_VMX\_EPT\_VPID\_CAP (see Appendix A.10) to determine whether the processor supports this feature.
3. N is the physical-address width supported by the logical processor. Software can determine a processor's physical-address width by executing CPUID with 80000008H in EAX. The physical-address width is returned in bits 7:0 of EAX.

The EPTP exists only on processors that support the 1-setting of the "enable EPT" VM-execution control.

### 24.6.12 Virtual-Processor Identifier (VPID)

The **virtual-processor identifier** (VPID) is a 16-bit field. It exists only on processors that support the 1-setting of the "enable VPID" VM-execution control. See Section 28.1 for details regarding the use of this field.

### 24.6.13 Controls for PAUSE-Loop Exiting

On processors that support the 1-setting of the "PAUSE-loop exiting" VM-execution control, the VM-execution control fields include the following 32-bit fields:

- **PLE\_Gap.** Software can configure this field as an upper bound on the amount of time between two successive executions of PAUSE in a loop.
- **PLE\_Window.** Software can configure this field as an upper bound on the amount of time a guest is allowed to execute in a PAUSE loop.

These fields measure time based on a counter that runs at the same rate as the timestamp counter (TSC). See Section 25.1.3 for more details regarding PAUSE-loop exiting.

### 24.6.14 VM-Function Controls

The **VM-function controls** constitute a 64-bit vector that governs use of the VMFUNC instruction in VMX non-root operation. This field is supported only on processors that support the 1-settings of both the “activate secondary controls” primary processor-based VM-execution control and the “enable VM functions” secondary processor-based VM-execution control.

Table 24-9 lists the VM-function controls. See Section 25.5.5 for more details of how these controls affect processor behavior in VMX non-root operation.

**Table 24-9. Definitions of VM-Function Controls**

| Bit Position(s) | Name           | Description  |
|-----------------|----------------|--|
| 0               | EPTP switching | The EPTP-switching VM function changes the EPT pointer to a value chosen from the EPTP list. See Section 25.5.5.3. |

All other bits in this field are reserved to 0. Software should consult the VMX capability MSR IA32\_VMX\_VMFUNC (see Appendix A.11) to determine which bits are reserved. Failure to clear reserved bits causes subsequent VM entries to fail (see Section 26.2.1.1).

Processors that support the 1-setting of the “EPTP switching” VM-function control also support a 64-bit field called the **EPTP-list address**. This field contains the physical address of the 4-KByte EPTP list. The EPTP list comprises 512 8-Byte entries (each an EPTP value) and is used by the EPTP-switching VM function (see Section 25.5.5.3).

### 24.6.15 VMCS Shadowing Bitmap Addresses

On processors that support the 1-setting of the “VMCS shadowing” VM-execution control, the VM-execution control fields include the 64-bit physical addresses of the **VMREAD bitmap** and the **VMWRITE bitmap**. Each bitmap is 4 KBytes in size and thus contains 32 KBits. The addresses are the **VMREAD-bitmap address** and the **VMWRITE-bitmap address**.

If the “VMCS shadowing” VM-execution control is 1, executions of VMREAD and VMWRITE may consult these bitmaps (see Section 24.10 and Section 30.3).

### 24.6.16 ENCLS-Exiting Bitmap

The **ENCLS-exiting bitmap** is a 64-bit field. If the “enable ENCLS exiting” VM-execution control is 1, execution of ENCLS causes a VM exit if the bit in this field corresponding to the value of EAX is 1. If the bit is 0, the instruction executes normally. See Section 25.1.3 for more information.

### 24.6.17 Control Field for Page-Modification Logging

The **PML address** is a 64-bit field. It is the 4-KByte aligned address of the **page-modification log**. The page-modification log consists of 512 64-bit entries. It is used for the page-modification logging feature. Details of the page-modification logging are given in Section 28.2.5.

If the “enable PML” VM-execution control is 1, VM entry ensures that the PML address is 4-KByte aligned. The PML address exists only on processors that support the 1-setting of the “enable PML” VM-execution control.

### 24.6.18 Controls for Virtualization Exceptions

On processors that support the 1-setting of the “EPT-violation #VE” VM-execution control, the VM-execution control fields include the following:

- **Virtualization-exception information address** (64 bits). This field contains the physical address of the **virtualization-exception information area**. When a logical processor encounters a virtualization exception, it saves virtualization-exception information at the virtualization-exception information address; see Section 25.5.6.2.

- **EPTP index** (16 bits). When an EPT violation causes a virtualization exception, the processor writes the value of this field to the virtualization-exception information area. The EPTP-switching VM function updates this field (see Section 25.5.5.3).

### 24.6.19 XSS-Exiting Bitmap

On processors that support the 1-setting of the “enable XSAVES/XRSTORS” VM-execution control, the VM-execution control fields include a 64-bit **XSS-exiting bitmap**. If the “enable XSAVES/XRSTORS” VM-execution control is 1, executions of XSAVES and XRSTORS may consult this bitmap (see Section 25.1.3 and Section 25.3).

## 24.7 VM-EXIT CONTROL FIELDS

The VM-exit control fields govern the behavior of VM exits. They are discussed in Section 24.7.1 and Section 24.7.2.

### 24.7.1 VM-Exit Controls

The **VM-exit controls** constitute a 32-bit vector that governs the basic operation of VM exits. Table 24-10 lists the controls supported. See Chapter 27 for complete details of how these controls affect VM exits.

**Table 24-10. Definitions of VM-Exit Controls**

| Bit Position(s) | Name                            | Description  |
|-----------------|---------------------------------|--|
| 2               | Save debug controls             | This control determines whether DR7 and the IA32_DEBUGCTL MSR are saved on VM exit. The first processors to support the virtual-machine extensions supported only the 1-setting of this control.   |
| 9               | Host address-space size         | On processors that support Intel 64 architecture, this control determines whether a logical processor is in 64-bit mode after the next VM exit. Its value is loaded into CS.L, IA32_EFER.LME, and IA32_EFER.LMA on every VM exit. <sup>1</sup><br>This control must be 0 on processors that do not support Intel 64 architecture.  |
| 12              | Load IA32_PERF_GLOBAL_CTRL      | This control determines whether the IA32_PERF_GLOBAL_CTRL MSR is loaded on VM exit.  |
| 15              | Acknowledge interrupt on exit   | This control affects VM exits due to external interrupts: <ul style="list-style-type: none"> <li>▪ If such a VM exit occurs and this control is 1, the logical processor acknowledges the interrupt controller, acquiring the interrupt’s vector. The vector is stored in the VM-exit interruption-information field, which is marked valid.</li> <li>▪ If such a VM exit occurs and this control is 0, the interrupt is not acknowledged and the VM-exit interruption-information field is marked invalid.</li> </ul> |
| 18              | Save IA32_PAT                   | This control determines whether the IA32_PAT MSR is saved on VM exit.  |
| 19              | Load IA32_PAT                   | This control determines whether the IA32_PAT MSR is loaded on VM exit.   |
| 20              | Save IA32_EFER                  | This control determines whether the IA32_EFER MSR is saved on VM exit.   |
| 21              | Load IA32_EFER                  | This control determines whether the IA32_EFER MSR is loaded on VM exit.  |
| 22              | Save VMX-preemption timer value | This control determines whether the value of the VMX-preemption timer is saved on VM exit.   |
| 23              | Clear IA32_BNDCFGS              | This control determines whether the IA32_BNDCFGS MSR is cleared on VM exit.  |
| 24              | Conceal VMX from PT             | If this control is 1, Intel Processor Trace does not produce a paging information packet (PIP) on a VM exit or a VMCS packet on an SMM VM exit (see Chapter 35).   |

**NOTES:**

1. Since Intel 64 architecture specifies that IA32\_EFER.LMA is always set to the logical-AND of CRO.PG and IA32\_EFER.LME, and since CRO.PG is always 1 in VMX operation, IA32\_EFER.LMA is always identical to IA32\_EFER.LME in VMX operation.

All other bits in this field are reserved, some to 0 and some to 1. Software should consult the VMX capability MSRs IA32\_VMX\_EXIT\_CTLS and IA32\_VMX\_TRUE\_EXIT\_CTLS (see Appendix A.4) to determine how it should set the reserved bits. Failure to set reserved bits properly causes subsequent VM entries to fail (see Section 26.2.1.2).

The first processors to support the virtual-machine extensions supported only the 1-settings of bits 0–8, 10, 11, 13, 14, 16, and 17. The VMX capability MSR IA32\_VMX\_EXIT\_CTLS always reports that these bits must be 1. Logical processors that support the 0-settings of any of these bits will support the VMX capability MSR IA32\_VMX\_TRUE\_EXIT\_CTLS MSR, and software should consult this MSR to discover support for the 0-settings of these bits. Software that is not aware of the functionality of any one of these bits should set that bit to 1.

### 24.7.2 VM-Exit Controls for MSRs

A VMM may specify lists of MSRs to be stored and loaded on VM exits. The following VM-exit control fields determine how MSRs are stored on VM exits:

- **VM-exit MSR-store count** (32 bits). This field specifies the number of MSRs to be stored on VM exit. It is recommended that this count not exceed 512 bytes.<sup>1</sup> Otherwise, unpredictable processor behavior (including a machine check) may result during VM exit.
- **VM-exit MSR-store address** (64 bits). This field contains the physical address of the VM-exit MSR-store area. The area is a table of entries, 16 bytes per entry, where the number of entries is given by the VM-exit MSR-store count. The format of each entry is given in Table 24-11. If the VM-exit MSR-store count is not zero, the address must be 16-byte aligned.

**Table 24-11. Format of an MSR Entry**

| Bit Position(s) | Contents  |
|-----------------|-----------|
| 31:0            | MSR index |
| 63:32           | Reserved  |
| 127:64          | MSR data  |

See Section 27.4 for how this area is used on VM exits.

The following VM-exit control fields determine how MSRs are loaded on VM exits:

- **VM-exit MSR-load count** (32 bits). This field contains the number of MSRs to be loaded on VM exit. It is recommended that this count not exceed 512 bytes. Otherwise, unpredictable processor behavior (including a machine check) may result during VM exit.<sup>2</sup>
- **VM-exit MSR-load address** (64 bits). This field contains the physical address of the VM-exit MSR-load area. The area is a table of entries, 16 bytes per entry, where the number of entries is given by the VM-exit MSR-load count (see Table 24-11). If the VM-exit MSR-load count is not zero, the address must be 16-byte aligned.

See Section 27.6 for how this area is used on VM exits.

## 24.8 VM-ENTRY CONTROL FIELDS

The VM-entry control fields govern the behavior of VM entries. They are discussed in Sections 24.8.1 through 24.8.3.

1. Future implementations may allow more MSRs to be stored reliably. Software should consult the VMX capability MSR IA32\_VMX\_MISC to determine the number supported (see Appendix A.6).
2. Future implementations may allow more MSRs to be loaded reliably. Software should consult the VMX capability MSR IA32\_VMX\_MISC to determine the number supported (see Appendix A.6).

## 24.8.1 VM-Entry Controls

The VM-entry controls constitute a 32-bit vector that governs the basic operation of VM entries. Table 24-12 lists the controls supported. See Chapter 24 for how these controls affect VM entries.

**Table 24-12. Definitions of VM-Entry Controls**

| Bit Position(s) | Name                              | Description   |
|-----------------|-----------------------------------|---|
| 2               | Load debug controls               | This control determines whether DR7 and the IA32_DEBUGCTL MSR are loaded on VM entry. The first processors to support the virtual-machine extensions supported only the 1-setting of this control.  |
| 9               | IA-32e mode guest                 | On processors that support Intel 64 architecture, this control determines whether the logical processor is in IA-32e mode after VM entry. Its value is loaded into IA32_EFER.LMA as part of VM entry. <sup>1</sup><br>This control must be 0 on processors that do not support Intel 64 architecture. |
| 10              | Entry to SMM                      | This control determines whether the logical processor is in system-management mode (SMM) after VM entry. This control must be 0 for any VM entry from outside SMM.  |
| 11              | Deactivate dual-monitor treatment | If set to 1, the default treatment of SMIs and SMM is in effect after the VM entry (see Section 34.15.7). This control must be 0 for any VM entry from outside SMM.   |
| 13              | Load IA32_PERF_GLOBAL_CTRL        | This control determines whether the IA32_PERF_GLOBAL_CTRL MSR is loaded on VM entry.  |
| 14              | Load IA32_PAT                     | This control determines whether the IA32_PAT MSR is loaded on VM entry.   |
| 15              | Load IA32_EFER                    | This control determines whether the IA32_EFER MSR is loaded on VM entry.  |
| 16              | Load IA32_BNDCFGS                 | This control determines whether the IA32_BNDCFGS MSR is loaded on VM entry.   |
| 17              | Conceal VMX from PT               | If this control is 1, Intel Processor Trace does not produce a paging information packet (PIP) on a VM entry or a VMCS packet on a VM entry that returns from SMM (see Chapter 35).   |

### NOTES:

1. Bit 5 of the IA32\_VMX\_MISC MSR is read as 1 on any logical processor that supports the 1-setting of the “unrestricted guest” VM-execution control. If it is read as 1, every VM exit stores the value of IA32\_EFER.LMA into the “IA-32e mode guest” VM-entry control (see Section 27.2).

All other bits in this field are reserved, some to 0 and some to 1. Software should consult the VMX capability MSRs IA32\_VMX\_ENTRY\_CTLS and IA32\_VMX\_TRUE\_ENTRY\_CTLS (see Appendix A.5) to determine how it should set the reserved bits. Failure to set reserved bits properly causes subsequent VM entries to fail (see Section 26.2.1.3).

The first processors to support the virtual-machine extensions supported only the 1-settings of bits 0–8 and 12. The VMX capability MSR IA32\_VMX\_ENTRY\_CTLS always reports that these bits must be 1. Logical processors that support the 0-settings of any of these bits will support the VMX capability MSR IA32\_VMX\_TRUE\_ENTRY\_CTLS MSR, and software should consult this MSR to discover support for the 0-settings of these bits. Software that is not aware of the functionality of any one of these bits should set that bit to 1.

## 24.8.2 VM-Entry Controls for MSRs

A VMM may specify a list of MSRs to be loaded on VM entries. The following VM-entry control fields manage this functionality:

- **VM-entry MSR-load count** (32 bits). This field contains the number of MSRs to be loaded on VM entry. It is recommended that this count not exceed 512 bytes. Otherwise, unpredictable processor behavior (including a machine check) may result during VM entry.<sup>1</sup>

1. Future implementations may allow more MSRs to be loaded reliably. Software should consult the VMX capability MSR IA32\_VMX\_MISC to determine the number supported (see Appendix A.6).

- **VM-entry MSR-load address** (64 bits). This field contains the physical address of the VM-entry MSR-load area. The area is a table of entries, 16 bytes per entry, where the number of entries is given by the VM-entry MSR-load count. The format of entries is described in Table 24-11. If the VM-entry MSR-load count is not zero, the address must be 16-byte aligned.

See Section 26.4 for details of how this area is used on VM entries.

### 24.8.3 VM-Entry Controls for Event Injection

VM entry can be configured to conclude by delivering an event through the IDT (after all guest state and MSRs have been loaded). This process is called **event injection** and is controlled by the following three VM-entry control fields:

- **VM-entry interruption-information field** (32 bits). This field provides details about the event to be injected. Table 24-13 describes the field.

**Table 24-13. Format of the VM-Entry Interruption-Information Field**

| Bit Position(s) | Content  |
|-----------------|--|
| 7:0             | Vector of interrupt or exception   |
| 10:8            | Interruption type:<br>0: External interrupt<br>1: Reserved<br>2: Non-maskable interrupt (NMI)<br>3: Hardware exception<br>4: Software interrupt<br>5: Privileged software exception<br>6: Software exception<br>7: Other event |
| 11              | Deliver error code (0 = do not deliver; 1 = deliver)   |
| 30:12           | Reserved   |
| 31              | Valid  |

- The **vector** (bits 7:0) determines which entry in the IDT is used or which other event is injected.
- The **interruption type** (bits 10:8) determines details of how the injection is performed. In general, a VMM should use the type **hardware exception** for all exceptions other than breakpoint exceptions (#BP; generated by INT3) and overflow exceptions (#OF; generated by INTO); it should use the type **software exception** for #BP and #OF. The type **other event** is used for injection of events that are not delivered through the IDT.
- For exceptions, the **deliver-error-code bit** (bit 11) determines whether delivery pushes an error code on the guest stack.
- VM entry injects an event if and only if the **valid bit** (bit 31) is 1. The valid bit in this field is cleared on every VM exit (see Section 27.2).
- **VM-entry exception error code** (32 bits). This field is used if and only if the valid bit (bit 31) and the deliver-error-code bit (bit 11) are both set in the VM-entry interruption-information field.
- **VM-entry instruction length** (32 bits). For injection of events whose type is software interrupt, software exception, or privileged software exception, this field is used to determine the value of RIP that is pushed on the stack.

See Section 26.5 for details regarding the mechanics of event injection, including the use of the interruption type and the VM-entry instruction length.

VM exits clear the valid bit (bit 31) in the VM-entry interruption-information field.



## 24.9 VM-EXIT INFORMATION FIELDS

The VMCS contains a section of fields that contain information about the most recent VM exit.

On some processors, attempts to write to these fields with VMWRITE fail (see “VMWRITE—Write Field to Virtual-Machine Control Structure” in Chapter 30).<sup>1</sup>

### 24.9.1 Basic VM-Exit Information

The following VM-exit information fields provide basic information about a VM exit:

- **Exit reason** (32 bits). This field encodes the reason for the VM exit and has the structure given in Table 24-14.

**Table 24-14. Format of Exit Reason**

| Bit Position(s) | Contents   |
|-----------------|--|
| 15:0            | Basic exit reason  |
| 26:16           | Reserved (cleared to 0)  |
| 27              | A VM exit saves this bit as 1 to indicate that the VM exit was incident to enclave mode. |
| 28              | Pending MTF VM exit  |
| 29              | VM exit from VMX root operation  |
| 30              | Reserved (cleared to 0)  |
| 31              | VM-entry failure (0 = true VM exit; 1 = VM-entry failure)                                |

- Bits 15:0 provide basic information about the cause of the VM exit (if bit 31 is clear) or of the VM-entry failure (if bit 31 is set). Appendix C enumerates the basic exit reasons.
- Bit 28 is set only by an SMM VM exit (see Section 34.15.2) that took priority over an MTF VM exit (see Section 25.5.2) that would have occurred had the SMM VM exit not occurred. See Section 34.15.2.3.
- Bit 29 is set if and only if the processor was in VMX root operation at the time the VM exit occurred. This can happen only for SMM VM exits. See Section 34.15.2.
- Because some VM-entry failures load processor state from the host-state area (see Section 26.7), software must be able to distinguish such cases from true VM exits. Bit 31 is used for that purpose.
- **Exit qualification** (64 bits; 32 bits on processors that do not support Intel 64 architecture). This field contains additional information about the cause of VM exits due to the following: debug exceptions; page-fault exceptions; start-up IPIs (SIPIs); task switches; INVEPT; INVLPG; INVVPID; LGDT; LIDT; LLDT; LTR; SGDT; SIDT; SLDT; STR; VMCLEAR; VMPTRLD; VMPTRST; VMREAD; VMWRITE; VMXON; XRSTORS; XSAVES; control-register accesses; MOV DR; I/O instructions; and MWAIT. The format of the field depends on the cause of the VM exit. See Section 27.2.1 for details.
- **Guest-linear address** (64 bits; 32 bits on processors that do not support Intel 64 architecture). This field is used in the following cases:
  - VM exits due to attempts to execute LMSW with a memory operand.
  - VM exits due to attempts to execute INS or OUTS.
  - VM exits due to system-management interrupts (SMIs) that arrive immediately after retirement of I/O instructions.
  - Certain VM exits due to EPT violations
 See Section 27.2.1 and Section 34.15.2.3 for details of when and how this field is used.

1. Software can discover whether these fields can be written by reading the VMX capability MSR IA32\_VMX\_MISC (see Appendix A.6).

- **Guest-physical address** (64 bits). This field is used VM exits due to EPT violations and EPT misconfigurations. See Section 27.2.1 for details of when and how this field is used.

### 24.9.2 Information for VM Exits Due to Vectored Events

Event-specific information is provided for VM exits due to the following vectored events: exceptions (including those generated by the instructions INT3, INTO, BOUND, and UD); external interrupts that occur while the “acknowledge interrupt on exit” VM-exit control is 1; and non-maskable interrupts (NMIs). This information is provided in the following fields:

- **VM-exit interruption information** (32 bits). This field receives basic information associated with the event causing the VM exit. Table 24-15 describes this field.

**Table 24-15. Format of the VM-Exit Interruption-Information Field**

| Bit Position(s) | Content   |
|-----------------|---|
| 7:0             | Vector of interrupt or exception  |
| 10:8            | Interruption type:<br>0: External interrupt<br>1: Not used<br>2: Non-maskable interrupt (NMI)<br>3: Hardware exception<br>4 - 5: Not used<br>6: Software exception<br>7: Not used |
| 11              | Error code valid (0 = invalid; 1 = valid)   |
| 12              | NMI unblocking due to IRET  |
| 30:13           | Reserved (cleared to 0)   |
| 31              | Valid   |

- **VM-exit interruption error code** (32 bits). For VM exits caused by hardware exceptions that would have delivered an error code on the stack, this field receives that error code.

Section 27.2.2 provides details of how these fields are saved on VM exits.

### 24.9.3 Information for VM Exits That Occur During Event Delivery

Additional information is provided for VM exits that occur during event delivery in VMX non-root operation.<sup>1</sup> This information is provided in the following fields:

- **IDT-vectoring information** (32 bits). This field receives basic information associated with the event that was being delivered when the VM exit occurred. Table 24-16 describes this field.

---

1. This includes cases in which the event delivery was caused by event injection as part of VM entry; see Section 26.5.1.2.

**Table 24-16. Format of the IDT-Vectoring Information Field**

| Bit Position(s) | Content   |
|-----------------|---|
| 7:0             | Vector of interrupt or exception  |
| 10:8            | Interruption type:<br>0: External interrupt<br>1: Not used<br>2: Non-maskable interrupt (NMI)<br>3: Hardware exception<br>4: Software interrupt<br>5: Privileged software exception<br>6: Software exception<br>7: Not used |
| 11              | Error code valid (0 = invalid; 1 = valid)   |
| 12              | Undefined   |
| 30:13           | Reserved (cleared to 0)   |
| 31              | Valid   |

- **IDT-vectoring error code** (32 bits). For VM exits that occur during delivery of hardware exceptions that would have delivered an error code on the stack, this field receives that error code.

See Section 27.2.3 provides details of how these fields are saved on VM exits.

## 24.9.4 Information for VM Exits Due to Instruction Execution

The following fields are used for VM exits caused by attempts to execute certain instructions in VMX non-root operation:

- **VM-exit instruction length** (32 bits). For VM exits resulting from instruction execution, this field receives the length in bytes of the instruction whose execution led to the VM exit.<sup>1</sup> See Section 27.2.4 for details of when and how this field is used.
- **VM-exit instruction information** (32 bits). This field is used for VM exits due to attempts to execute `INS`, `INVEPT`, `INVVPID`, `LIDT`, `LGDT`, `LLDT`, `LTR`, `OUTS`, `SIDT`, `SGDT`, `SLDT`, `STR`, `VMCLEAR`, `VMPTRLD`, `VMPTRST`, `VMREAD`, `VMWRITE`, or `VMXON`.<sup>2</sup> The format of the field depends on the cause of the VM exit. See Section 27.2.4 for details.

The following fields (64 bits each; 32 bits on processors that do not support Intel 64 architecture) are used only for VM exits due to SMIs that arrive immediately after retirement of I/O instructions. They provide information about that I/O instruction:

- **I/O RCX**. The value of RCX before the I/O instruction started.
- **I/O RSI**. The value of RSI before the I/O instruction started.
- **I/O RDI**. The value of RDI before the I/O instruction started.
- **I/O RIP**. The value of RIP before the I/O instruction started (the RIP that addressed the I/O instruction).

## 24.9.5 VM-Instruction Error Field

The 32-bit **VM-instruction error field** does not provide information about the most recent VM exit. In fact, it is not modified on VM exits. Instead, it provides information about errors encountered by a non-faulting execution of one of the VMX instructions.

1. This field is also used for VM exits that occur during the delivery of a software interrupt or software exception.

2. Whether the processor provides this information on VM exits due to attempts to execute `INS` or `OUTS` can be determined by consulting the VMX capability MSR `IA32_VMX_BASIC` (see Appendix A.1).

## 24.10 VMCS TYPES: ORDINARY AND SHADOW

Every VMCS is either an **ordinary VMCS** or a **shadow VMCS**. A VMCS's type is determined by the shadow-VMCS indicator in the VMCS region (this is the value of bit 31 of the first 4 bytes of the VMCS region; see Table 24-1): 0 indicates an ordinary VMCS, while 1 indicates a shadow VMCS. Shadow VMCSs are supported only on processors that support the 1-setting of the "VMCS shadowing" VM-execution control (see Section 24.6.2).

A shadow VMCS differs from an ordinary VMCS in two ways:

- An ordinary VMCS can be used for VM entry but a shadow VMCS cannot. Attempts to perform VM entry when the current VMCS is a shadow VMCS fail (see Section 26.1).
- The VMREAD and VMWRITE instructions can be used in VMX non-root operation to access a shadow VMCS but not an ordinary VMCS. This fact results from the following:
  - If the "VMCS shadowing" VM-execution control is 0, execution of the VMREAD and VMWRITE instructions in VMX non-root operation always cause VM exits (see Section 25.1.3).
  - If the "VMCS shadowing" VM-execution control is 1, execution of the VMREAD and VMWRITE instructions in VMX non-root operation can access the VMCS referenced by the VMCS link pointer (see Section 30.3).
  - If the "VMCS shadowing" VM-execution control is 1, VM entry ensures that any VMCS referenced by the VMCS link pointer is a shadow VMCS (see Section 26.3.1.5).

In VMX root operation, both types of VMCSs can be accessed with the VMREAD and VMWRITE instructions.

Software should not modify the shadow-VMCS indicator in the VMCS region of a VMCS that is active. Doing so may cause the VMCS to become corrupted (see Section 24.11.1). Before modifying the shadow-VMCS indicator, software should execute VMCLEAR for the VMCS to ensure that it is not active.

## 24.11 SOFTWARE USE OF THE VMCS AND RELATED STRUCTURES

This section details guidelines that software should observe when using a VMCS and related structures. It also provides descriptions of consequences for failing to follow guidelines.

### 24.11.1 Software Use of Virtual-Machine Control Structures

To ensure proper processor behavior, software should observe certain guidelines when using an active VMCS.

No VMCS should ever be active on more than one logical processor. If a VMCS is to be "migrated" from one logical processor to another, the first logical processor should execute VMCLEAR for the VMCS (to make it inactive on that logical processor and to ensure that all VMCS data are in memory) before the other logical processor executes VMPTRLD for the VMCS (to make it active on the second logical processor).<sup>1</sup> A VMCS that is made active on more than one logical processor may become **corrupted** (see below).

Software should not modify the shadow-VMCS indicator (see Table 24-1) in the VMCS region of a VMCS that is active. Doing so may cause the VMCS to become corrupted. Before modifying the shadow-VMCS indicator, software should execute VMCLEAR for the VMCS to ensure that it is not active.

Software should use the VMREAD and VMWRITE instructions to access the different fields in the current VMCS (see Section 24.11.2). Software should never access or modify the VMCS data of an active VMCS using ordinary memory operations, in part because the format used to store the VMCS data is implementation-specific and not architecturally defined, and also because a logical processor may maintain some VMCS data of an active VMCS on the processor and not in the VMCS region. The following items detail some of the hazards of accessing VMCS data using ordinary memory operations:

- Any data read from a VMCS with an ordinary memory read does not reliably reflect the state of the VMCS. Results may vary from time to time or from logical processor to logical processor.

---

1. As noted in Section 24.1, execution of the VMPTRLD instruction makes a VMCS is active. In addition, VM entry makes active any shadow VMCS referenced by the VMCS link pointer in the current VMCS. If a shadow VMCS is made active by VM entry, it is necessary to execute VMCLEAR for that VMCS before allowing that VMCS to become active on another logical processor.

- Writing to a VMCS with an ordinary memory write is not guaranteed to have a deterministic effect on the VMCS. Doing so may cause the VMCS to become corrupted (see below).

(Software can avoid these hazards by removing any linear-address mappings to a VMCS region before executing a VMPTRLD for that region and by not remapping it until after executing VMCLEAR for that region.)

If a logical processor leaves VMX operation, any VMCSs active on that logical processor may be corrupted (see below). To prevent such corruption of a VMCS that may be used either after a return to VMX operation or on another logical processor, software should execute VMCLEAR for that VMCS before executing the VMXOFF instruction or removing power from the processor (e.g., as part of a transition to the S3 and S4 power states).

This section has identified operations that may cause a VMCS to become corrupted. These operations may cause the VMCS’s data to become undefined. Behavior may be unpredictable if that VMCS used subsequently on any logical processor. The following items detail some hazards of VMCS corruption:

- VM entries may fail for unexplained reasons or may load undesired processor state.
- The processor may not correctly support VMX non-root operation as documented in Chapter 27 and may generate unexpected VM exits.
- VM exits may load undesired processor state, save incorrect state into the VMCS, or cause the logical processor to transition to a shutdown state.

### 24.11.2 VMREAD, VMWRITE, and Encodings of VMCS Fields

Every field of the VMCS is associated with a 32-bit value that is its **encoding**. The encoding is provided in an operand to VMREAD and VMWRITE when software wishes to read or write that field. These instructions fail if given, in 64-bit mode, an operand that sets an encoding bit beyond bit 32. See Chapter 30 for a description of these instructions.

The structure of the 32-bit encodings of the VMCS components is determined principally by the width of the fields and their function in the VMCS. See Table 24-17.

**Table 24-17. Structure of VMCS Component Encoding**

| Bit Position(s) | Contents  |
|-----------------|---|
| 0               | Access type (0 = full; 1 = high); must be full for 16-bit, 32-bit, and natural-width fields |
| 9:1             | Index   |
| 11:10           | Type:<br>0: control<br>1: VM-exit information<br>2: guest state<br>3: host state            |
| 12              | Reserved (must be 0)  |
| 14:13           | Width:<br>0: 16-bit<br>1: 64-bit<br>2: 32-bit<br>3: natural-width                           |
| 31:15           | Reserved (must be 0)  |

The following items detail the meaning of the bits in each encoding:

- **Field width.** Bits 14:13 encode the width of the field.
  - A value of 0 indicates a 16-bit field.
  - A value of 1 indicates a 64-bit field.

- A value of 2 indicates a 32-bit field.
- A value of 3 indicates a **natural-width** field. Such fields have 64 bits on processors that support Intel 64 architecture and 32 bits on processors that do not.

Fields whose encodings use value 1 are specially treated to allow 32-bit software access to all 64 bits of the field. Such access is allowed by defining, for each such field, an encoding that allows direct access to the high 32 bits of the field. See below.

- **Field type.** Bits 11:10 encode the type of VMCS field: control, guest-state, host-state, or VM-exit information. (The last category also includes the VM-instruction error field.)
- **Index.** Bits 9:1 distinguish components with the same field width and type.
- **Access type.** Bit 0 must be 0 for all fields except for 64-bit fields (those with field-width 1; see above). A VMREAD or VMWRITE using an encoding with this bit cleared to 0 accesses the entire field. For a 64-bit field with field-width 1, a VMREAD or VMWRITE using an encoding with this bit set to 1 accesses only the high 32 bits of the field.

Appendix B gives the encodings of all fields in the VMCS.

The following describes the operation of VMREAD and VMWRITE based on processor mode, VMCS-field width, and access type:

- 16-bit fields:
  - A VMREAD returns the value of the field in bits 15:0 of the destination operand; other bits of the destination operand are cleared to 0.
  - A VMWRITE writes the value of bits 15:0 of the source operand into the VMCS field; other bits of the source operand are not used.
- 32-bit fields:
  - A VMREAD returns the value of the field in bits 31:0 of the destination operand; in 64-bit mode, bits 63:32 of the destination operand are cleared to 0.
  - A VMWRITE writes the value of bits 31:0 of the source operand into the VMCS field; in 64-bit mode, bits 63:32 of the source operand are not used.
- 64-bit fields and natural-width fields using the full access type outside IA-32e mode.
  - A VMREAD returns the value of bits 31:0 of the field in its destination operand; bits 63:32 of the field are ignored.
  - A VMWRITE writes the value of its source operand to bits 31:0 of the field and clears bits 63:32 of the field.
- 64-bit fields and natural-width fields using the full access type in 64-bit mode (only on processors that support Intel 64 architecture).
  - A VMREAD returns the value of the field in bits 63:0 of the destination operand
  - A VMWRITE writes the value of bits 63:0 of the source operand into the VMCS field.
- 64-bit fields using the high access type.
  - A VMREAD returns the value of bits 63:32 of the field in bits 31:0 of the destination operand; in 64-bit mode, bits 63:32 of the destination operand are cleared to 0.
  - A VMWRITE writes the value of bits 31:0 of the source operand to bits 63:32 of the field; in 64-bit mode, bits 63:32 of the source operand are not used.

Software seeking to read a 64-bit field outside IA-32e mode can use VMREAD with the full access type (reading bits 31:0 of the field) and VMREAD with the high access type (reading bits 63:32 of the field); the order of the two VMREAD executions is not important. Software seeking to modify a 64-bit field outside IA-32e mode should first use VMWRITE with the full access type (establishing bits 31:0 of the field while clearing bits 63:32) and then use VMWRITE with the high access type (establishing bits 63:32 of the field).

### 24.11.3 Initializing a VMCS

Software should initialize fields in a VMCS (using VMWRITE) before using the VMCS for VM entry. Failure to do so may result in unpredictable behavior; for example, a VM entry may fail for unexplained reasons, or a successful transition (VM entry or VM exit) may load processor state with unexpected values.

It is not necessary to initialize fields that the logical processor will not use. (For example, it is not necessary to initialize the MSR-bitmap address if the “use MSR bitmaps” VM-execution control is 0.)

A processor maintains some VMCS information that cannot be modified with the VMWRITE instruction; this includes a VMCS’s launch state (see Section 24.1). Such information may be stored in the VMCS data portion of a VMCS region. Because the format of this information is implementation-specific, there is no way for software to know, when it first allocates a region of memory for use as a VMCS region, how the processor will determine this information from the contents of the memory region.

In addition to its other functions, the VMCLEAR instruction initializes any implementation-specific information in the VMCS region referenced by its operand. To avoid the uncertainties of implementation-specific behavior, software should execute VMCLEAR on a VMCS region before making the corresponding VMCS active with VMPTRLD for the first time. (Figure 24-1 illustrates how execution of VMCLEAR puts a VMCS into a well-defined state.)

The following software usage is consistent with these limitations:

- VMCLEAR should be executed for a VMCS before it is used for VM entry for the first time.
- VMLAUNCH should be used for the first VM entry using a VMCS after VMCLEAR has been executed for that VMCS.
- VMRESUME should be used for any subsequent VM entry using a VMCS (until the next execution of VMCLEAR for the VMCS).

It is expected that, in general, VMRESUME will have lower latency than VMLAUNCH. Since “migrating” a VMCS from one logical processor to another requires use of VMCLEAR (see Section 24.11.1), which sets the launch state of the VMCS to “clear”, such migration requires the next VM entry to be performed using VMLAUNCH. Software developers can avoid the performance cost of increased VM-entry latency by avoiding unnecessary migration of a VMCS from one logical processor to another.

### 24.11.4 Software Access to Related Structures

In addition to data in the VMCS region itself, VMX non-root operation can be controlled by data structures that are referenced by pointers in a VMCS (for example, the I/O bitmaps). While the pointers to these data structures are parts of the VMCS, the data structures themselves are not. They are not accessible using VMREAD and VMWRITE but by ordinary memory writes.

Software should ensure that each such data structure is modified only when no logical processor with a current VMCS that references it is in VMX non-root operation. Doing otherwise may lead to unpredictable behavior (including behaviors identified in Section 24.11.1).

### 24.11.5 VMXON Region

Before executing VMXON, software allocates a region of memory (called the VMXON region)<sup>1</sup> that the logical processor uses to support VMX operation. The physical address of this region (the VMXON pointer) is provided in an operand to VMXON. The VMXON pointer is subject to the limitations that apply to VMCS pointers:

- The VMXON pointer must be 4-KByte aligned (bits 11:0 must be zero).
- The VMXON pointer must not set any bits beyond the processor’s physical-address width.<sup>2,3</sup>

1. The amount of memory required for the VMXON region is the same as that required for a VMCS region. This size is implementation specific and can be determined by consulting the VMX capability MSR IA32\_VMX\_BASIC (see Appendix A.1).

2. Software can determine a processor’s physical-address width by executing CPUID with 80000008H in EAX. The physical-address width is returned in bits 7:0 of EAX.

3. If IA32\_VMX\_BASIC[48] is read as 1, the VMXON pointer must not set any bits in the range 63:32; see Appendix A.1.

Before executing VMXON, software should write the VMCS revision identifier (see Section 24.2) to the VMXON region. (Specifically, it should write the 31-bit VMCS revision identifier to bits 30:0 of the first 4 bytes of the VMXON region; bit 31 should be cleared to 0.) It need not initialize the VMXON region in any other way. Software should use a separate region for each logical processor and should not access or modify the VMXON region of a logical processor between execution of VMXON and VMXOFF on that logical processor. Doing otherwise may lead to unpredictable behavior (including behaviors identified in Section 24.11.1).



## 14. Updates to Chapter 35, Volume 3C

Change bars show changes to Chapter 35 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C: System Programming Guide, Part 3*.

-----  
Changes to this chapter: Information on PMI Preservation added, and typo corrections.

## CHAPTER 35

# INTEL® PROCESSOR TRACE

---

## 35.1 OVERVIEW

Intel® Processor Trace (Intel PT) is an extension of Intel® Architecture that captures information about software execution using dedicated hardware facilities that cause only minimal performance perturbation to the software being traced. This information is collected in data packets. The initial implementations of Intel PT offer **control flow tracing**, which generates a variety of packets to be processed by a software decoder. The packets include timing, program flow information (e.g. branch targets, branch taken/not taken indications) and program-induced mode related information (e.g. Intel TSX state transitions, CR3 changes). These packets may be buffered internally before being sent to the memory subsystem or other output mechanism available in the platform. Debug software can process the trace data and reconstruct the program flow.

Later generations include additional trace sources, including software trace instrumentation using PTWRITE, and Power Event tracing.

### 35.1.1 Features and Capabilities

Intel PT's control flow trace generates a variety of packets that, when combined with the binaries of a program by a post-processing tool, can be used to produce an exact execution trace. The packets record flow information such as instruction pointers (IP), indirect branch targets, and directions of conditional branches within contiguous code regions (basic blocks).

Intel PT can also be configured to log software-generated packets using PTWRITE, and packets describing processor power management events.

In addition, the packets record other contextual, timing, and bookkeeping information that enables both functional and performance debugging of applications. Intel PT has several control and filtering capabilities available to customize the tracing information collected and to append other processor state and timing information to enable debugging. For example, there are modes that allow packets to be filtered based on the current privilege level (CPL) or the value of CR3.

Configuration of the packet generation and filtering capabilities are programmed via a set of MSRs. The MSRs generally follow the naming convention of IA32\_RTIT\_\*. The capability provided by these configuration MSRs are enumerated by CPUID, see Section 35.3. Details of the MSRs for configuring Intel PT are described in Section 35.2.7.

#### 35.1.1.1 Packet Summary

After a tracing tool has enabled and configured the appropriate MSRs, the processor will collect and generate trace information in the following categories of packets (for more details on the packets, see Section 35.4):

- Packets about basic information on program execution; these include:
  - Packet Stream Boundary (PSB) packets: PSB packets act as 'heartbeats' that are generated at regular intervals (e.g., every 4K trace packet bytes). These packets allow the packet decoder to find the packet boundaries within the output data stream; a PSB packet should be the first packet that a decoder looks for when beginning to decode a trace.
  - Paging Information Packet (PIP): PIPs record modifications made to the CR3 register. This information, along with information from the operating system on the CR3 value of each process, allows the debugger to attribute linear addresses to their correct application source.
  - Time-Stamp Counter (TSC) packets: TSC packets aid in tracking wall-clock time, and contain some portion of the software-visible time-stamp counter.
  - Core Bus Ratio (CBR) packets: CBR packets contain the core:bus clock ratio.

- Overflow (OVF) packets: OVF packets are sent when the processor experiences an internal buffer overflow, resulting in packets being dropped. This packet notifies the decoder of the loss and can help the decoder to respond to this situation.
- Packets about control flow information:
  - Taken Not-Taken (TNT) packets: TNT packets track the “direction” of direct conditional branches (taken or not taken).
  - Target IP (TIP) packets: TIP packets record the target IP of indirect branches, exceptions, interrupts, and other branches or events. These packets can contain the IP, although that IP value may be compressed by eliminating upper bytes that match the last IP. There are various types of TIP packets; they are covered in more detail in Section 35.4.2.2.
  - Flow Update Packets (FUP): FUPs provide the source IP addresses for asynchronous events (interrupt and exceptions), as well as other cases where the source address cannot be determined from the binary.
  - **MODE** packets: These packets provide the decoder with important processor execution information so that it can properly interpret the dis-assembled binary and trace log. MODE packets have a variety of formats that indicate details such as the execution mode (16-bit, 32-bit, or 64-bit).
- Packets inserted by software:
  - PTWRITE (PTW) packets: includes the value of the operand passed to the PTWRITE instruction (see “PTWRITE - Write Data to a Processor Trace Packet” in *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B*).
- Packets about processor power management events:
  - MWAIT packets: Indicate successful completion of an MWAIT operation to a C-state deeper than C0.0.
  - Power State Entry (PWRE) packets: Indicate entry to a C-state deeper than C0.0.
  - Power State Exit (PWRX) packets: Indicate exit from a C-state deeper than C0.0, returning to C0.
  - Execution Stopped (EXSTOP) packets: Indicate that software execution has stopped, due to events such as P-state change, C-state change, or thermal throttling.

## 35.2 INTEL® PROCESSOR TRACE OPERATIONAL MODEL

This section describes the overall Intel Processor Trace mechanism and the essential concepts relevant to how it operates.

### 35.2.1 Change of Flow Instruction (COFI) Tracing

A basic program block is a section of code where no jumps or branches occur. The instruction pointers (IPs) in this block of code need not be traced, as the processor will execute them from start to end without redirecting code flow. Instructions such as branches, and events such as exceptions or interrupts, can change the program flow. These instructions and events that change program flow are called Change of Flow Instructions (COFI). There are three categories of COFI:

- Direct transfer COFI.
- Indirect transfer COFI.
- Far transfer COFI.

The following subsections describe the COFI events that result in trace packet generation. Table 35-1 lists branch instruction by COFI types. For detailed description of specific instructions, see *Intel® 64 and IA-32 Architectures Software Developer’s Manual*.

**Table 35-1. COFI Type for Branch Instructions**

| COFI Type          | Instructions   |
|--------------------|--|
| Conditional Branch | JA, JAE, JB, JBE, JC, JCXZ, JECXZ, JRCXZ, JE, JG, JGE, JL, JLE, JNA, JNAE, JNB, JNBE, JNC, JNE, JNG, JNGE, JNL, JNLE, JNO, JNP, JNS, JNZ, JO, JP, JPE, JPO, JS, JZ, LOOP, LOOPE, LOOPNE, LOOPNZ, LOOPZ |

Table 35-1. COFI Type for Branch Instructions

| COFI Type                   | Instructions  |
|-----------------------------|---|
| Unconditional Direct Branch | JMP (E9 xx, EB xx), CALL (E8 xx)  |
| Indirect Branch             | JMP (FF /4), CALL (FF /2)   |
| Near Ret                    | RET (C3, C2 xx)   |
| Far Transfers               | INT3, INTn, INTO, IRET, IRETD, IRETQ, JMP (EA xx, FF /5), CALL (9A xx, FF /3), RET (CB, CA xx), SYS-CALL, SYSRET, SYSENTER, SYSEXIT, VMLAUNCH, VMRESUME |

### 35.2.1.1 Direct Transfer COFI

Direct Transfer COFI are relative branches. This means that their target is an IP whose offset from the current IP is embedded in the instruction bytes. It is not necessary to indicate target of these instructions in the trace output since it can be obtained through the source disassembly. Conditional branches need to indicate only whether the branch is taken or not. Unconditional branches do not need any recording in the trace output. There are two sub-categories:

- **Conditional Branch (Jcc, J\*CXZ) and LOOP**

To track this type of instruction, the processor encodes a single bit (taken or not taken — TNT) to indicate the program flow after the instruction.

Jcc, J\*CXZ, and LOOP can be traced with TNT bits. To improve the trace packet output efficiency, the processor will compact several TNT bits into a single packet.

- **Unconditional Direct Jumps**

There is no trace output required for direct unconditional jumps (like JMP near relative or CALL near relative) since they can be directly inferred from the application assembly. Direct unconditional jumps do not generate a TNT bit or a Target IP packet, though TIP.PGD and TIP.PGE packets can be generated by unconditional direct jumps that toggle Intel PT enables (see Section 35.2.5).

### 35.2.1.2 Indirect Transfer COFI

Indirect transfer instructions involve updating the IP from a register or memory location. Since the register or memory contents can vary at any time during execution, there is no way to know the target of the indirect transfer until the register or memory contents are read. As a result, the disassembled code is not sufficient to determine the target of this type of COFI. Therefore, tracing hardware must send out the destination IP in the trace packet for debug software to determine the target address of the COFI. Note that this IP may be a linear or effective address (see Section 35.3.1.1).

An indirect transfer instruction generates a Target IP Packet (TIP) that contains the target address of the branch. There are two sub-categories:

- **Near JMP Indirect and Near Call Indirect**

As previously mentioned, the target of an indirect COFI resides in the contents of either a register or memory location. Therefore, the processor must generate a packet that includes this target address to allow the decoder to determine the program flow.

- **Near RET**

When a CALL instruction executes, it pushes onto the stack the address of the next instruction following the CALL. Upon completion of the call procedure, the RET instruction is often used to pop the return address off of the call stack and redirect code flow back to the instruction following the CALL.

A RET instruction simply transfers program flow to the address it popped off the stack. Because a called procedure may change the return address on the stack before executing the RET instruction, debug software can be misled if it assumes that code flow will return to the instruction following the last CALL. Therefore, even for near RET, a Target IP Packet may be sent.

- **RET Compression**

A special case is applied if the target of the RET is consistent with what would be expected from tracking the CALL stack. If it is assured that the decoder has seen the corresponding CALL (with “corresponding” defined

as the CALL with matching stack depth), and the RET target is the instruction after that CALL, the RET target may be “compressed”. In this case, only a single TNT bit of “taken” is generated instead of a Target IP Packet. To ensure that the decoder will not be confused in cases of RET compression, only RETs that correspond to CALLs which have been seen since the last PSB packet may be compressed in a given logical processor. For details, see “Indirect Transfer Compression for Returns (RET)” in Section 35.4.2.2.

### 35.2.1.3 Far Transfer COFI

All operations that change the instruction pointer and are not near jumps are “far transfers”. This includes exceptions, interrupts, traps, TSX aborts, and instructions that do far transfers.

All far transfers will produce a Target IP (TIP) packet, which provides the destination IP address. For those far transfers that cannot be inferred from the binary source (e.g., asynchronous events such as exceptions and interrupts), the TIP will be preceded by a Flow Update packet (FUP), which provides the source IP address at which the event was taken. Table 35-23 indicates exactly which IP will be included in the FUP generated by a far transfer.

## 35.2.2 Software Trace Instrumentation with PTWRITE

PTWRITE provides a mechanism by which software can instrument the Intel PT trace. PTWRITE is a ring3-accessible instruction that can be passed to a register or memory variable, see “PTWRITE - Write Data to a Processor Trace Packet” in *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B* for details. The contents of that variable will be used as the payload for the PTW packet (see Table 35-40 “PTW Packet Definition”), inserted at the time of PTWRITE retirement, assuming PTWRITE is enabled and all other filtering conditions are met. Decode and analysis software will then be able to determine the meaning of the PTWRITE packet based on the IP of the associated PTWRITE instruction.

PTWRITE is enabled via IA32\_RTIT\_CTL.PTWEn[12] (see Table 35-6). Optionally, the user can use IA32\_RTIT\_CTL.FUPonPTW[5] to enable PTW packets to be followed by FUP packets containing the IP of the associated PTWRITE instruction. Support for PTWRITE is introduced in Intel® Atom™ processors based on the Goldmont Plus microarchitecture.

## 35.2.3 Power Event Tracing

Power Event Trace is a capability that exposes core- and thread-level sleep state and power down transition information. When this capability is enabled, the trace will expose information about:

- Scenarios where software execution stops.
  - Due to sleep state entry, frequency change, or other powerdown.
  - Includes the IP, when in the tracing context.
- The requested and resolved hardware thread C-state.
  - Including indication of hardware autonomous C-state entry.
- The last and deepest core C-state achieved during a sleep session.
- The reason for C-state wake.

This information is in addition to the bus ratio (CBR) information provided by default after any powerdown, and the timing information (TSC, TMA, MTC, CYC) provided during or after a powerdown state.

Power Event Trace is enabled via IA32\_RTIT\_CTL.PwrEvtEn[4]. Support for Power Event Tracing is introduced in Intel® Atom™ processors based on the Goldmont Plus microarchitecture.

## 35.2.4 Trace Filtering

Intel Processor Trace provides filtering capabilities, by which the debug/profile tool can control what code is traced.

### 35.2.4.1 Filtering by Current Privilege Level (CPL)

Intel PT provides the ability to configure a logical processor to generate trace packets only when CPL = 0, when CPL > 0, or regardless of CPL.

CPL filtering ensures that no IPs or other architectural state information associated with the filtered CPL can be seen in the log. For example, if the processor is configured to trace only when CPL > 0, and software executes SYSCALL (changing the CPL to 0), the destination IP of the SYSCALL will be suppressed from the generated packet (see the discussion of TIP.PGD in Section 35.4.2.5).

It should be noted that CPL is always 0 in real-address mode and that CPL is always 3 in virtual-8086 mode. To trace code in these modes, filtering should be configured accordingly.

When software is executing in a non-enabled CPL, ContextEn is cleared. See Section 35.2.5.1 for details.

### 35.2.4.2 Filtering by CR3

Intel PT supports a CR3-filtering mechanism by which the generation of packets containing architectural states can be enabled or disabled based on the value of CR3. A debugger can use CR3 filtering to trace only a single application without context switching the state of the RTIT MSRs. For the reconstruction of traces from software with multiple threads, debug software may wish to context-switch for the state of the RTIT MSRs (if the operating system does not provide context-switch support) to separate the output for the different threads (see Section 35.3.5, "Context Switch Consideration").

To trace for only a single CR3 value, software can write that value to the IA32\_RTIT\_CR3\_MATCH MSR, and set IA32\_RTIT\_CTL.CR3Filter. When CR3 value does not match IA32\_RTIT\_CR3\_MATCH and IA32\_RTIT\_CTL.CR3Filter is 1, ContextEn is forced to 0, and packets containing architectural states will not be generated. Some other packets can be generated when ContextEn is 0; see Section 35.2.5.3 for details. When CR3 does match IA32\_RTIT\_CR3\_MATCH (or when IA32\_RTIT\_CTL.CR3Filter is 0), CR3 filtering does not force ContextEn to 0 (although it could be 0 due to other filters or modes).

CR3 matches IA32\_RTIT\_CR3\_MATCH if the two registers are identical for bits 63:12, or 63:5 when in PAE paging mode; the lower 5 bits of CR3 and IA32\_RTIT\_CR3\_MATCH are ignored. CR3 filtering is independent of the value of CR0.PG.

When CR3 filtering is in use, PIP packets may still be seen in the log if the processor is configured to trace when CPL = 0 (IA32\_RTIT\_CTL.OS = 1). If not, no PIP packets will be seen.

### 35.2.4.3 Filtering by IP

Trace packet generation with configurable filtering by IP is supported if CPUID.(EAX=14H, ECX=0):EBX[bit 2] = 1. Intel PT can be configured to enable the generation of packets containing architectural states only when the processor is executing code within certain IP ranges. If the IP is outside of these ranges, generation of some packets is blocked.

IP filtering is enabled using the ADDRn\_CFG fields in the IA32\_RTIT\_CTL MSR (Section 35.2.7.2), where the digit 'n' is a zero-based number that selects which address range is being configured. Each ADDRn\_CFG field configures the use of the register pair IA32\_RTIT\_ADDRn\_A and IA32\_RTIT\_ADDRn\_B (Section 35.2.7.5). IA32\_RTIT\_ADDRn\_A defines the base and IA32\_RTIT\_ADDRn\_B specifies the limit of the range in which tracing is enabled. Thus each range, referred to as the ADDRn range, is defined by [IA32\_RTIT\_ADDRn\_A, IA32\_RTIT\_ADDRn\_B]. There can be multiple such ranges, software can query CPUID (Section 35.3.1) for the number of ranges supported on a processor.

Default behavior (ADDRn\_CFG=0) defines no IP filter range, meaning FilterEn is always set. In this case code at any IP can be traced, though other filters, such as CR3 or CPL, could limit tracing. When ADDRn\_CFG is set to enable IP filtering (see Section 35.3.1), tracing will commence when a taken branch or event is seen whose target address is in the ADDRn range.

While inside a tracing region and with FilterEn is set, leaving the tracing region may only be detected once a taken branch or event with a target outside the range is retired. If an ADDRn range is entered or exited by executing the next sequential instruction, rather than by a control flow transfer, FilterEn may not toggle immediately. See Section 35.2.5.5 for more details on FilterEn.

Note that these address range base and limit values are inclusive, such that the range includes the first and last instruction whose first instruction byte is in the ADDRn range.

Depending upon processor implementation, IP filtering may be based on linear or effective address. This can cause different behavior between implementations if CSbase is not equal to zero or in real mode. See Section 35.3.1.1 for details. Software can query CPUID to determine filters are based on linear or effective address (Section 35.3.1).

Note that some packets, such as MTC (Section 35.3.7) and other timing packets, do not depend on FilterEn. For details on which packets depend on FilterEn, and hence are impacted by IP filtering, see Section 35.4.1.

### TraceStop

The ADDRn ranges can also be configured to cause tracing to be disabled upon entry to the specified region. This is intended for cases where unexpected code is executed, and the user wishes to immediately stop generating packets in order to avoid overwriting previously written packets.

The TraceStop mechanism works much the same way that IP filtering does, and uses the same address comparison logic. The TraceStop region base and limit values are programmed into one or more ADDRn ranges, but IA32\_RTIT\_CTL.ADDRn\_CFG is configured with the TraceStop encoding. Like FilterEn, TraceStop is detected when a taken branch or event lands in a TraceStop region.

Further, TraceStop requires that TriggerEn=1 at the beginning of the branch/event, and ContextEn=1 upon completion of the branch/event. When this happens, the CPU will set IA32\_RTIT\_STATUS.Stopped, thereby clearing TriggerEn and hence disabling packet generation. This may generate a TIP.PGD packet with the target IP of the branch or event that entered the TraceStop region. Finally, a TraceStop packet will be inserted, to indicate that the condition was hit.

If a TraceStop condition is encountered during buffer overflow (Section 35.3.8), it will not be dropped, but will instead be signaled once the overflow has resolved.

Note that a TraceStop event does not guarantee that all internally buffered packets are flushed out of internal buffers. To ensure that this has occurred, the user should clear TraceEn.

To resume tracing after a TraceStop event, the user must first disable Intel PT by clearing IA32\_RTIT\_CTL.TraceEn before the IA32\_RTIT\_STATUS.Stopped bit can be cleared. At this point Intel PT can be reconfigured, and tracing resumed.

Note that the IA32\_RTIT\_STATUS.Stopped bit can also be set using the ToPA STOP bit. See Section 35.2.6.2.

### IP Filtering Example

The following table gives an example of IP filtering behavior. Assume that IA32\_RTIT\_ADDRn\_A = the IP of RangeBase, and that IA32\_RTIT\_ADDRn\_B = the IP of RangeLimit, while IA32\_RTIT\_CTL.ADDRn\_CFG = 0x1 (enable ADDRn range as a FilterEn range).

**Table 35-2. IP Filtering Packet Example**

| Code Flow  | Packets  |
|--|--|
| <pre> Bar:     jmp RangeBase // jump into filter range RangeBase:     jcc Foo // not taken     add eax, 1 Foo:     jmp RangeLimit+1 // jump out of filter range RangeLimit:     nop     jcc Bar                     </pre> | <pre> TIP.PGE(RangeBase) TNT(0) TIP.PGD(RangeLimit+1)                     </pre> |

### IP Filtering and TraceStop

It is possible for the user to configure IP filter range(s) and TraceStop range(s) that overlap. In this case, code executing in the non-overlapping portion of either range will behave as would be expected from that range. Code executing in the overlapping range will get TraceStop behavior.

## 35.2.5 Packet Generation Enable Controls

Intel Processor Trace includes a variety of controls that determine whether a packet is generated. In general, most packets are sent only if Packet Enable (**PacketEn**) is set. **PacketEn** is an internal state maintained in hardware in response to software configurable enable controls, **PacketEn** is not visible to software directly. The relationship of **PacketEn** to the software-visible controls in the configuration MSRs is described in this section.

### 35.2.5.1 Packet Enable (**PacketEn**)

When **PacketEn** is set, the processor is in the mode that Intel PT is monitoring and all packets can be generated to log what is being executed. **PacketEn** is composed of other states according to this relationship:

$$\text{PacketEn} \leftarrow \text{TriggerEn} \text{ AND } \text{ContextEn} \text{ AND } \text{FilterEn} \text{ AND } \text{BranchEn}$$

These constituent controls are detailed in the following subsections.

**PacketEn** ultimately determines when the processor is tracing. When **PacketEn** is set, all control flow packets are enabled. When **PacketEn** is clear, no control flow packets are generated, though other packets (timing and book-keeping packets) may still be sent. See Section 35.2.6 for details of **PacketEn** and packet generation.

Note that, on processors that do not support IP filtering (i.e., `CPUID.(EAX=14H, ECX=0):EBX.IPFILT_WRSTPRSV[bit 2] = 0`), **FilterEn** is treated as always set.

### 35.2.5.2 Trigger Enable (**TriggerEn**)

Trigger Enable (**TriggerEn**) is the primary indicator that trace packet generation is active. **TriggerEn** is set when `IA32_RTIT_CTL.TraceEn` is set, and cleared by any of the following conditions:

- `TraceEn` is cleared by software.
- A `TraceStop` condition is encountered and `IA32_RTIT_STATUS.Stopped` is set.
- `IA32_RTIT_STATUS.Error` is set due to an operational error (see Section 35.3.9).

Software can discover the current **TriggerEn** value by reading the `IA32_RTIT_STATUS.TriggerEn` bit. When **TriggerEn** is clear, tracing is inactive and no packets are generated.

### 35.2.5.3 Context Enable (**ContextEn**)

Context Enable (**ContextEn**) indicates whether the processor is in the state or mode that software configured hardware to trace. For example, if execution with `CPL = 0` code is not being traced (`IA32_RTIT_CTL.OS = 0`), then **ContextEn** will be 0 when the processor is in `CPL0`.

Software can discover the current **ContextEn** value by reading the `IA32_RTIT_STATUS.ContextEn` bit. **ContextEn** is defined as follows:

$$\begin{aligned} \text{ContextEn} = & !((\text{IA32\_RTIT\_CTL.OS} = 0 \text{ AND } \text{CPL} = 0) \text{ OR} \\ & (\text{IA32\_RTIT\_CTL.USER} = 0 \text{ AND } \text{CPL} > 0) \text{ OR } (\text{IS\_IN\_A\_PRODUCTION\_ENCLAVE}^1) \text{ OR} \\ & (\text{IA32\_RTIT\_CTL.CR3Filter} = 1 \text{ AND } \text{IA32\_RTIT\_CR3\_MATCH} \text{ does not match CR3}) \end{aligned}$$

If the clearing of **ContextEn** causes **PacketEn** to be cleared, a Packet Generation Disable (`TIP.PGD`) packet is generated, but its IP payload is suppressed. If the setting of **ContextEn** causes **PacketEn** to be set, a Packet Generation Enable (`TIP.PGE`) packet is generated.

When **ContextEn** is 0, control flow packets (`TNT`, `FUP`, `TIP.*`, `MODE.*`) are not generated, and no Linear Instruction Pointers (LIPs) are exposed. However, some packets, such as `MTC` and `PSB` (see Section 35.4.2.16 and Section 35.4.2.17), may still be generated while **ContextEn** is 0. For details of which packets are generated only when **ContextEn** is set, see Section 35.4.1.

The processor does not update **ContextEn** when **TriggerEn** = 0.

The value of **ContextEn** will toggle only when **TriggerEn** = 1.

1. Trace packets generation is disabled in a production enclave, see Section 35.2.8.5. See *Intel® Software Guard Extensions Programming Reference* about differences between a production enclave and a debug enclave.



### 35.2.5.4 Branch Enable (BranchEn)

This value is based purely on the IA32\_RTIT\_CTL.BranchEn value. If **BranchEn** is not set, then relevant COFI packets (TNT, TIP\*, FUP, MODE.\*) are suppressed. Other packets related to timing (TSC, TMA, MTC, CYC), as well as PSB, will be generated normally regardless. Further, PIP and VMCS continue to be generated, as indicators of what software is running.

### 35.2.5.5 Filter Enable (FilterEn)

Filter Enable indicates that the Instruction Pointer (IP) is within the range of IPs that Intel PT is configured to watch. Software can get the state of Filter Enable by a RDMSR of IA32\_RTIT\_STATUS.FilterEn. For details on configuration and use of IP filtering, see Section 35.2.4.3.

On clearing of FilterEn that also clears PacketEn, a Packet Generation Disable (TIP.PGD) will be generated, but unlike the ContextEn case, the IP payload may not be suppressed. For direct, unconditional branches, as well as for indirect branches (including RETs), the PGD generated by leaving the tracing region and clearing FilterEn will contain the target IP. This means that IPs from outside the configured range can be exposed in the trace, as long as they are within context.

When FilterEn is 0, control flow packets are not generated (e.g., TNT, TIP). However, some packets, such as PIP, MTC, and PSB, may still be generated while FilterEn is clear. For details on packet enable dependencies, see Section 35.4.1.

After TraceEn is set, FilterEn is set to 1 at all times if there is no IP filter range configured by software (IA32\_RTIT\_CTL.ADDRn\_CFG != 1, for all n), or if the processor does not support IP filtering (i.e., CPUID.(EAX=14H, ECX=0):EBX.IPFILT\_WRSTPRSV[bit 2] = 0). FilterEn will toggle only when TraceEn=1 and ContextEn=1, and when at least one range is configured for IP filtering.

## 35.2.6 Trace Output

Intel PT output should be viewed independently from trace content and filtering mechanisms. The options available for trace output can vary across processor generations and platforms.

Trace output is written out using one of the following output schemes, as configured by the ToPA and FabricEn bit fields of IA32\_RTIT\_CTL (see Section 35.2.7.2):

- A single, contiguous region of physical address space.
- A collection of variable-sized regions of physical memory. These regions are linked together by tables of pointers to those regions, referred to as Table of Physical Addresses (**ToPA**). The trace output stores bypass the caches and the TLBs, but are not serializing. This is intended to minimize the performance impact of the output.
- A platform-specific trace transport subsystem.

Regardless of the output scheme chosen, Intel PT stores bypass the processor caches by default. This ensures that they don't consume precious cache space, but they do not have the serializing aspects associated with un-cacheable (UC) stores. Software should avoid using MTRRs to mark any portion of the Intel PT output region as UC, as this may override the behavior described above and force Intel PT stores to UC, thereby incurring severe performance impact.

There is no guarantee that a packet will be written to memory or other trace endpoint after some fixed number of cycles after a packet-producing instruction executes. The only way to assure that all packets generated have reached their endpoint is to clear TraceEn and follow that with a store, fence, or serializing instruction; doing so ensures that all buffered packets are flushed out of the processor.

### 35.2.6.1 Single Range Output

When IA32\_RTIT\_CTL.ToPA and IA32\_RTIT\_CTL.FabricEn bits are clear, trace packet output is sent to a single, contiguous memory (or MMIO if DRAM is not available) range defined by a base address in IA32\_RTIT\_OUTPUT\_BASE (Section 35.2.7.7) and mask value in IA32\_RTIT\_OUTPUT\_MASK\_PTRS (Section 35.2.7.8). The current write pointer in this range is also stored in IA32\_RTIT\_OUTPUT\_MASK\_PTRS. This output range is circular, meaning that when the writes wrap around the end of the buffer they begin again at the base address.

This output method is best suited for cases where Intel PT output is either:

- Configured to be directed to a sufficiently large contiguous region of DRAM.
- Configured to go to an MMIO debug port, in order to route Intel PT output to a platform-specific trace endpoint (e.g., JTAG). In this scenario, a specific range of addresses is written in a circular manner, and SoC will intercept these writes and direct them to the proper device. Repeated writes to the same address do not overwrite each other, but are accumulated by the debugger, and hence no data is lost by the circular nature of the buffer.

The processor will determine the address to which to write the next trace packet output byte as follows:

```
OutputBase[63:0] ← IA32_RTIT_OUTPUT_BASE[63:0]
OutputMask[63:0] ← ZeroExtend64(IA32_RTIT_OUTPUT_MASK_PTRS[31:0])
OutputOffset[63:0] ← ZeroExtend64(IA32_RTIT_OUTPUT_MASK_PTRS[63:32])
trace_store_phys_addr ← (OutputBase & ~OutputMask) + (OutputOffset & OutputMask)
```

### Single-Range Output Errors

If the output base and mask are not properly configured by software, an operational error (see Section 35.3.9) will be signaled, and tracing disabled. Error scenarios with single-range output are:

- Mask value is non-contiguous.  
IA32\_RTIT\_OUTPUT\_MASK\_PTRS.MaskOrTablePointer value has a 0 in a less significant bit position than the most significant bit containing a 1.
- Base address and Mask are mis-aligned, and have overlapping bits set.  
IA32\_RTIT\_OUTPUT\_BASE && IA32\_RTIT\_OUTPUT\_MASK\_PTRS.MaskOrTableOffset > 0.
- Illegal Output Offset  
IA32\_RTIT\_OUTPUT\_MASK\_PTRS.OutputOffset is greater than the mask value (IA32\_RTIT\_OUTPUT\_MASK\_PTRS.MaskOrTableOffset).

Also note that errors can be signaled due to trace packet output overlapping with restricted memory, see Section 35.2.6.4.

### 35.2.6.2 Table of Physical Addresses (ToPA)

When IA32\_RTIT\_CTL.ToPA is set and IA32\_RTIT\_CTL.FabricEn is clear, the ToPA output mechanism is utilized. The ToPA mechanism uses a linked list of tables; see Figure 35-1 for an illustrative example. Each entry in the table contains some attribute bits, a pointer to an output region, and the size of the region. The last entry in the table may hold a pointer to the next table. This pointer can either point to the top of the current table (for circular array) or to the base of another table. The table size is not fixed, since the link to the next table can exist at any entry.

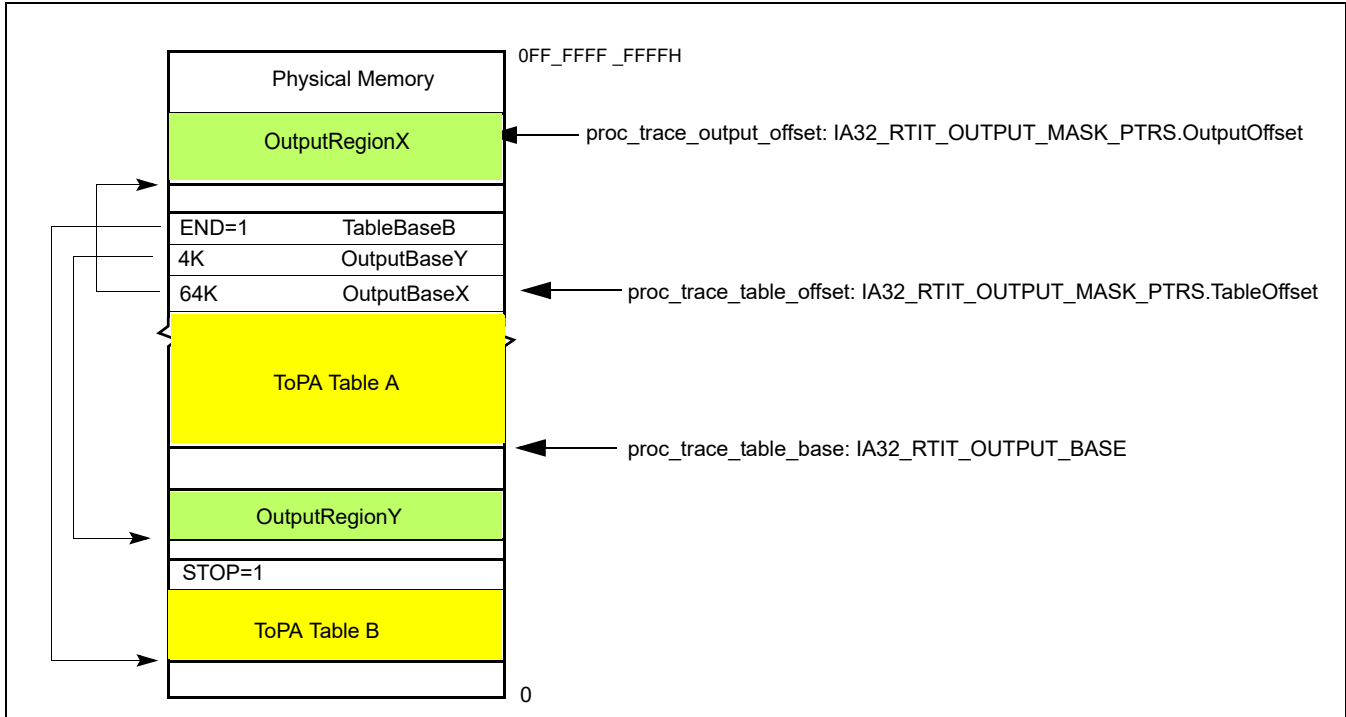
The processor treats the various output regions referenced by the ToPA table(s) as a unified buffer. This means that a single packet may span the boundary between one output region and the next.

The ToPA mechanism is controlled by three values maintained by the processor:

- **proc\_trace\_table\_base.**  
This is the physical address of the base of the current ToPA table. When tracing is enabled, the processor loads this value from the IA32\_RTIT\_OUTPUT\_BASE MSR. While tracing is enabled, the processor updates the IA32\_RTIT\_OUTPUT\_BASE MSR with changes to proc\_trace\_table\_base, but these updates may not be synchronous to software execution. When tracing is disabled, the processor ensures that the MSR contains the latest value of proc\_trace\_table\_base.
- **proc\_trace\_table\_offset.**  
This indicates the entry of the current table that is currently in use. (This entry contains the address of the current output region.) When tracing is enabled, the processor loads this value from bits 31:7 (MaskOrTableOffset) of the IA32\_RTIT\_OUTPUT\_MASK\_PTRS. While tracing is enabled, the processor updates IA32\_RTIT\_OUTPUT\_MASK\_PTRS.MaskOrTableOffset with changes to proc\_trace\_table\_offset, but these updates may not be synchronous to software execution. When tracing is disabled, the processor ensures that the MSR contains the latest value of proc\_trace\_table\_offset.
- **proc\_trace\_output\_offset.**  
This is a pointer into the current output region and indicates the location of the next write. When tracing is

enabled, the processor loads this value from bits 63:32 (OutputOffset) of the IA32\_RTIT\_OUTPUT\_MASK\_PTRS. While tracing is enabled, the processor updates IA32\_RTIT\_OUTPUT\_MASK\_PTRS.OutputOffset with changes to proc\_trace\_output\_offset, but these updates may not be synchronous to software execution. When tracing is disabled, the processor ensures that the MSR contains the latest value of proc\_trace\_output\_offset.

Figure 35-1 provides an illustration (not to scale) of the table and associated pointers.



**Figure 35-1. ToPA Memory Illustration**

With the ToPA mechanism, the processor writes packets to the current output region (identified by `proc_trace_table_base` and the `proc_trace_table_offset`). The offset within that region to which the next byte will be written is identified by `proc_trace_output_offset`. When that region is filled with packet output (thus `proc_trace_output_offset = RegionSize-1`), `proc_trace_table_offset` is moved to the next ToPA entry, `proc_trace_output_offset` is set to 0, and packet writes begin filling the new output region specified by `proc_trace_table_offset`.

As packets are written out, each store derives its physical address as follows:

$$\text{trace\_store\_phys\_addr} \leftarrow \text{Base address from current ToPA table entry} + \text{proc\_trace\_output\_offset}$$

Eventually, the regions represented by all entries in the table may become full, and the final entry of the table is reached. An entry can be identified as the final entry because it has either the END or STOP attribute. The END attribute indicates that the address in the entry does not point to another output region, but rather to another ToPA table. The STOP attribute indicates that tracing will be disabled once the corresponding region is filled. See Table 35-3 and the section that follows for details on STOP.

When an END entry is reached, the processor loads `proc_trace_table_base` with the base address held in this END entry, thereby moving the current table pointer to this new table. The `proc_trace_table_offset` is reset to 0, as is the `proc_trace_output_offset`, and packet writes will resume at the base address indicated in the first entry.

If the table has no STOP or END entry, and trace-packet generation remains enabled, eventually the maximum table size will be reached (`proc_trace_table_offset = 01FFFFFFH`). In this case, the `proc_trace_table_offset` and `proc_trace_output_offset` are reset to 0 (wrapping back to the beginning of the current table) once the last output region is filled.

It is important to note that processor updates to the IA32\_RTIT\_OUTPUT\_BASE and IA32\_RTIT\_OUTPUT\_MASK\_PTRS MSRs are asynchronous to instruction execution. Thus, reads of these MSRs while Intel PT is enabled may return stale values. Like all IA32\_RTIT\_\* MSRs, the values of these MSRs should not be trusted or saved unless trace packet generation is first disabled by clearing IA32\_RTIT\_CTL.TraceEn. This ensures that the output MSR values account for all packets generated to that point, after which the processor will cease updating the output MSR values until tracing resumes.<sup>1</sup>

The processor may cache internally any number of entries from the current table or from tables that it references (directly or indirectly). If tracing is enabled, the processor may ignore or delay detection of modifications to these tables. To ensure that table changes are detected by the processor in a predictable manner, software should clear TraceEn before modifying the current table (or tables that it references) and only then re-enable packet generation.

### Single Output Region ToPA Implementation

The first processor generation to implement Intel PT supports only ToPA configurations with a single ToPA entry followed by an END entry that points back to the first entry (creating one circular output buffer). Such processors enumerate CPUID.(EAX=14H,ECX=0):ECX.MENTRY[bit 1] = 0 and CPUID.(EAX=14H,ECX=0):ECX.TOPAOUT[bit 0] = 1.

If CPUID.(EAX=14H,ECX=0):ECX.MENTRY[bit 1] = 0, ToPA tables can hold only one output entry, which must be followed by an END=1 entry which points back to the base of the table. Hence only one contiguous block can be used as output.

The lone output entry can have INT or STOP set, but nonetheless must be followed by an END entry as described above. Note that, if INT=1, the PMI will actually be delivered before the region is filled.

### ToPA Table Entry Format

The format of ToPA table entries is shown in Figure 35-2. The size of the address field is determined by the processor's physical-address width (MAXPHYADDR) in bits, as reported in CPUID.80000008H:EAX[7:0].

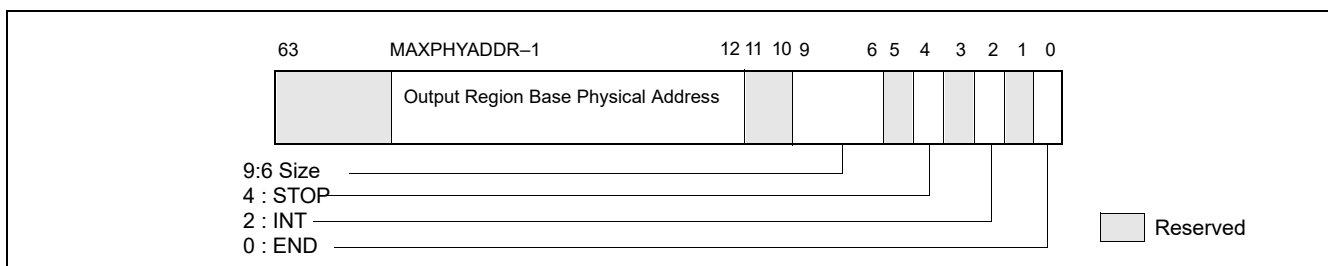


Figure 35-2. Layout of ToPA Table Entry

Table 35-3 describes the details of the ToPA table entry fields. If reserved bits are set to 1, an error is signaled.

Table 35-3. ToPA Table Entry Fields

| ToPA Entry Field                    | Description  |
|-------------------------------------|--|
| Output Region Base Physical Address | If END=0, this is the base physical address of the output region specified by this entry. Note that all regions must be aligned based on their size. Thus a 2M region must have bits 20:12 clear. If the region is not properly aligned, an operational error will be signaled when the entry is reached.<br>If END=1, this is the 4K-aligned base physical address of the next ToPA table (which may be the base of the current table, or the first table in the linked list if a circular buffer is desired). If the processor supports only a single ToPA output region (see above), this address must be the value currently in the IA32_RTIT_OUTPUT_BASE MSR. |

1. Although WRMSR is a serializing instruction, the execution of WRMSR that forces packet writes by clearing TraceEn does not itself cause these writes to be globally observed.

**Table 35-3. ToPA Table Entry Fields (Contd.)**

| ToPA Entry Field | Description   |
|------------------|---|
| Size             | Indicates the size of the associated output region. Encodings are:<br>0: 4K, 1: 8K, 2: 16K, 3: 32K, 4: 64K, 5: 128K, 6: 256K, 7: 512K,<br>8: 1M, 9: 2M, 10: 4M, 11: 8M, 12: 16M, 13: 32M, 14: 64M, 15: 128M<br>This field is ignored if END=1.  |
| STOP             | When the output region indicated by this entry is filled, software should disable packet generation. This will be accomplished by setting IA32_RTIT_STATUS.Stopped, which clears TriggerEn. This bit must be 0 if END=1; otherwise it is treated as reserved bit violation (see ToPA Errors).   |
| INT              | When the output region indicated by this entry is filled, signal Perfmon LVT interrupt.<br>Note that if both INT and STOP are set in the same entry, the STOP will happen before the INT. Thus the interrupt handler should expect that the IA32_RTIT_STATUS.Stopped bit will be set, and will need to be reset before tracing can be resumed.<br>This bit must be 0 if END=1; otherwise it is treated as reserved bit violation (see ToPA Errors). |
| END              | If set, indicates that this is an END entry, and thus the address field points to a table base rather than an output region base.<br>If END=1, INT and STOP must be set to 0; otherwise it is treated as reserved bit violation (see ToPA Errors). The Size field is ignored in this case.<br>If the processor supports only a single ToPA output region (see above), END must be set in the second table entry.                                    |

### ToPA STOP

Each ToPA entry has a STOP bit. If this bit is set, the processor will set the IA32\_RTIT\_STATUS.Stopped bit when the corresponding trace output region is filled. This will clear TriggerEn and thereby cease packet generation. See Section 35.2.7.4 for details on IA32\_RTIT\_STATUS.Stopped. This sequence is known as “ToPA Stop”.

No TIP.PGD packet will be seen in the output when the ToPA stop occurs, since the disable happens only when the region is already full. When this occurs, output ceases after the last byte of the region is filled, which may mean that a packet is cut off in the middle. Any packets remaining in internal buffers are lost and cannot be recovered.

When ToPA stop occurs, the IA32\_RTIT\_OUTPUT\_BASE MSR will hold the base address of the table whose entry had STOP=1. IA32\_RTIT\_OUTPUT\_MASK\_PTRS.MaskOffsetTableOffset will hold the index value for that entry, and the IA32\_RTIT\_OUTPUT\_MASK\_PTRS.OutputOffset should be set to the size of the region.

Note that this means the offset pointer is pointing to the next byte after the end of the region, a configuration that would produce an operational error if the configuration remained when tracing is re-enabled with IA32\_RTIT\_STATUS.Stopped cleared.

### ToPA PMI

Each ToPA entry has an INT bit. If this bit is set, the processor will signal a performance-monitoring interrupt (PMI) when the corresponding trace output region is filled. This interrupt is not precise, and it is thus likely that writes to the next region will occur by the time the interrupt is taken.

The following steps should be taken to configure this interrupt:

1. Enable PMI via the LVT Performance Monitor register (at MMIO offset 340H in xAPIC mode; via MSR 834H in x2APIC mode). See *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3B* for more details on this register. For ToPA PMI, set all fields to 0, save for the interrupt vector, which can be selected by software.
2. Set up an interrupt handler to service the interrupt vector that a ToPA PMI can raise.
3. Set the interrupt flag by executing STI.
4. Set the INT bit in the ToPA entry of interest and enable packet generation, using the ToPA output option. Thus, TraceEn=ToPA=1 in the IA32\_RTIT\_CTL MSR.

Once the INT region has been filled with packet output data, the interrupt will be signaled. This PMI can be distinguished from others by checking bit 55 (Trace\_ToPA\_PMI) of the IA32\_PERF\_GLOBAL\_STATUS MSR (MSR 38EH). Once the ToPA PMI handler has serviced the relevant buffer, writing 1 to bit 55 of the MSR at 390H (IA32\_GLOBAL\_STATUS\_RESET) clears IA32\_PERF\_GLOBAL\_STATUS.Trace\_ToPA\_PMI.

Intel PT is not frozen on PMI, and thus the interrupt handler will be traced (though filtering can prevent this). The Freeze\_Perfmon\_on\_PMI and Freeze\_LBRs\_on\_PMI settings in IA32\_DEBUGCTL will be applied on ToPA PMI just as on other PMIs, and hence Perfmon counters are frozen.

Assuming the PMI handler wishes to read any buffered packets for persistent output, or wishes to modify any Intel PT MSRs, software should first disable packet generation by clearing TraceEn. This ensures that all buffered packets are written to memory and avoids tracing of the PMI handler. The configuration MSRs can then be used to determine where tracing has stopped. If packet generation is disabled by the handler, it should then be manually re-enabled before the IRET if continued tracing is desired.

In rare cases, it may be possible to trigger a second ToPA PMI before the first is handled. This can happen if another ToPA region with INT=1 is filled before, or shortly after, the first PMI is taken, perhaps due to EFLAGS.IF being cleared for an extended period of time. This can manifest in two ways: either the second PMI is triggered before the first is taken, and hence only one PMI is taken, or the second is triggered after the first is taken, and thus will be taken when the handler for the first completes. Software can minimize the likelihood of the second case by clearing TraceEn at the beginning of the PMI handler. Further, it can detect such cases by then checking the Interrupt Request Register (IRR) for PMI pending, and checking the ToPA table base and off-set pointers (in IA32\_RTIT\_OUTPUT\_BASE and IA32\_RTIT\_OUTPUT\_MASK\_PTRS) to see if multiple entries with INT=1 have been filled.

When IA32\_RTIT\_CTL.InjectPsbPmiOnEnable[56] = 1, the PMI handler should take the following actions:

1. Ignore ToPA PMIs that are taken when TraceEn = 0, because the Intel PT MSR state may have already been saved by XSAVES, and because the PMI will be re-injected when Intel PT is re-enabled.
2. Clear the new IA32\_RTIT\_STATUS.PendTopaPMI[7] bit once the PMI has been handled. This bit should not be cleared in cases where a PMI is ignored due to TraceEn = 0.

### ToPA PMI and Single Output Region ToPA Implementation

A processor that supports only a single ToPA output region implementation (such that only one output region is supported; see above) will attempt to signal a ToPA PMI interrupt before the output wraps and overwrites the top of the buffer. To support this functionality, the PMI handler should disable packet generation as soon as possible.

Due to PMI skid, it is possible that, in rare cases, the wrap will have occurred before the PMI is delivered. Software can avoid this by setting the STOP bit in the ToPA entry (see Table 35-3); this will disable tracing once the region is filled, and no wrap will occur. This approach has the downside of disabling packet generation so that some of the instructions that led up to the PMI will not be traced. If the PMI skid is significant enough to cause the region to fill and tracing to be disabled, the PMI handler will need to clear the IA32\_RTIT\_STATUS.Stopped indication before tracing can resume.

### ToPA PMI and XSAVES/XRSTORS State Handling

In some cases the ToPA PMI may be taken after completion of an XSAVES instruction that switches Intel PT state, and in such cases any modification of Intel PT MSRs within the PMI handler will not persist when the saved Intel PT context is later restored with XRSTORS. To account for such a scenario, it is recommended that the Intel PT output configuration be modified by altering the ToPA tables themselves, rather than the Intel PT output MSRs. On processors that support PMI preservation (CPUID.(EAX=14H, ECX=0):EBX.INJECTPSBPMI[6] = 1), setting IA32\_RTIT\_CTL.InjectPsbPmiOnEnable[56] = 1 will ensure that a PMI that is pending at the time PT is disabled will be recorded by setting IA32\_RTIT\_STATUS.PendTopaPMI[7] = 1. A PMI will then be pended when the saved PT context is later restored.

Table 35-4 depicts a recommended PMI handler algorithm for managing multi-region ToPA output and handling ToPA PMIs that may arrive between XSAVES and XRSTORS. This algorithm is flexible to allow software to choose between adding entries to the current ToPA table, adding a new ToPA table, or using the current ToPA table as a circular buffer. It assumes that the ToPA entry that triggers the PMI is not the last entry in the table, which is the recommended treatment.

**Table 35-4. Algorithm to Manage Intel PT ToPA PMI and XSAVES/XRSTORS**

| Pseudo Code Flow   |
|--|
| <pre> IF (IA32_PERF_GLOBAL_STATUS.ToPA)   Save IA32_RTIT_CTL value;   IF ( IA32_RTIT_CTL.TraceEN )     Disable Intel PT by clearing TraceEn;   FI;   IF ( there is space available to grow the current ToPA table )     Add one or more ToPA entries after the last entry in the ToPA table;     Point new ToPA entry address field(s) to new output region base(s);   ELSE     Modify an upcoming ToPA entry in the current table to have END=1;     IF (output should transition to a new ToPA table )       Point the address of the "END=1" entry of the current table to the new table base;     ELSE       /* Continue to use the current ToPA table, make a circular. */       Point the address of the "END=1" entry to the base of the current table;       Modify the ToPA entry address fields for filled output regions to point to new, unused output regions;       /* Filled regions are those with index in the range of 0 to (IA32_RTIT_MASK_PTRS.MaskOrTableOffset -1). */     FI;   FI;   Restore saved IA32_RTIT_CTL.value; FI; </pre> |

### ToPA Errors

When a malformed ToPA entry is found, an operational error results (see Section 35.3.9). A malformed entry can be any of the following:

1. **ToPA entry reserved bit violation.**  
This describes cases where a bit marked as reserved in Section 35.2.6.2 above is set to 1.
2. **ToPA alignment violation.**  
This includes cases where illegal ToPA entry base address bits are set to 1:
  - a. ToPA table base address is not 4KB-aligned. The table base can be from a WRMSR to IA32\_RTIT\_OUTPUT\_BASE, or from a ToPA entry with END=1.
  - b. ToPA entry base address is not aligned to the ToPA entry size (e.g., a 2MB region with base address[20:12] not equal to 0).
  - c. ToPA entry base address sets upper physical address bits not supported by the processor.
3. **Illegal ToPA Output Offset** (if IA32\_RTIT\_STATUS.Stopped=0).  
IA32\_RTIT\_OUTPUT\_MASK\_PTRS.OutputOffset is greater than or equal to the size of the current ToPA output region size.
4. **ToPA rules violations.**  
These are similar to ToPA entry reserved bit violations; they are cases when a ToPA entry is encountered with illegal field combinations. They include the following:
  - a. Setting the STOP or INT bit on an entry with END=1.
  - b. Setting the END bit in entry 0 of a ToPA table.
  - c. On processors that support only a single ToPA entry (see above), two additional illegal settings apply:
    - i) ToPA table entry 1 with END=0.
    - ii) ToPA table entry 1 with base address not matching the table base.

In all cases, the error will be logged by setting `IA32_RTIT_STATUS.Error`, thereby disabling tracing when the problematic ToPA entry is reached (when `proc_trace_table_offset` points to the entry containing the error). Any packet bytes that are internally buffered when the error is detected may be lost.

Note that operational errors may also be signaled due to attempts to access restricted memory. See Section 35.2.6.4 for details.

A tracing software have a range of flexibility using ToPA to manage the interaction of Intel PT with application buffers, see Section 35.5.

### 35.2.6.3 Trace Transport Subsystem

When `IA32_RTIT_CTL.FabricEn` is set, the `IA32_RTIT_CTL.ToPA` bit is ignored, and trace output is written to the trace transport subsystem. The endpoints of this transport are platform-specific, and details of configuration options should refer to the specific platform documentation. The `FabricEn` bit is available to be set if `CPUID(EAX=14H,ECX=0):EBX[bit 3] = 1`.

### 35.2.6.4 Restricted Memory Access

Packet output cannot be directed to any regions of memory that are restricted by the platform. In particular, all memory accesses on behalf of packet output are checked against the SMRR regions. If there is any overlap with these regions, trace data collection will not function properly. Exact processor behavior is implementation-dependent; Table 35-5 summarizes several scenarios.

**Table 35-5. Behavior on Restricted Memory Access**

| Scenario                              | Description   |
|---------------------------------------|---|
| ToPA output region overlaps with SMRR | Stores to the restricted memory region will be dropped, and that packet data will be lost. Any attempt to read from that restricted region will return all 1s. The processor also may signal an error (Section 35.3.9) and disable tracing when the output pointer reaches the restricted region. If packet generation remains enabled, then packet output may continue once stores are no longer directed to restricted memory (on wrap, or if the output region is larger than the restricted memory region). |
| ToPA table overlaps with SMRR         | The processor will signal an error (Section 35.3.9) and disable tracing when the ToPA write pointer ( <code>IA32_RTIT_OUTPUT_BASE + (proc_trace_table_offset &lt;&lt; 3)</code> ) enters the restricted region.   |

It should also be noted that packet output should not be routed to the 4KB APIC MMIO region, as defined by the `IA32_APIC_BASE` MSR. For details about the APIC, refer to *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*. No error is signaled for this case.

### Modifications to Restricted Memory Regions

It is recommended that software disable packet generation before modifying the SMRRs to change the scope of the SMRR regions. This is because the processor reserves the right to cache any number of ToPA table entries internally, after checking them against restricted memory ranges. Once cached, the entries will not be checked again, meaning one could potentially route packet output to a newly restricted region. Software can ensure that any cached entries are written to memory by clearing `IA32_RTIT_CTL.TraceEn`.

## 35.2.7 Enabling and Configuration MSRs

### 35.2.7.1 General Considerations

Trace packet generation is enabled and configured by a collection of model-specific registers (MSRs), which are detailed below. Some notes on the configuration MSR behavior:

- If Intel Processor Trace is not supported by the processor (see Section 35.3.1), RDMSR or WRMSR of the `IA32_RTIT_*` MSRs will cause `#GP`.
- A WRMSR to any of these configuration MSRs that begins and ends with `IA32_RTIT_CTL.TraceEn` set will `#GP` fault. Packet generation must be disabled before the configuration MSRs can be changed.



Note: Software may write the same value back to IA32\_RTIT\_CTL without #GP, even if TraceEn=1.

- All configuration MSRs for Intel PT are duplicated per logical processor
- For each configuration MSR, any MSR write that attempts to change bits marked reserved, or utilize encodings marked reserved, will cause a #GP fault.
- All configuration MSRs for Intel PT are cleared on a cold RESET.
  - If CPUID.(EAX=14H, ECX=0):EBX.IPFILT\_WRSTPRSV[bit 2] = 1, only the TraceEn bit is cleared on warm RESET; though this may have the impact of clearing other bits in IA32\_RTIT\_STATUS. Other MSR values of the trace configuration MSRs are preserved on warm RESET.
- The semantics of MSR writes to trace configuration MSRs in this chapter generally apply to explicit WRMSR to these registers, using VM-exit or VM-entry MSR load list to these MSRs, XRSTORS with requested feature bit map including XSAVE map component of state\_8 (corresponding to IA32\_XSS[bit 8]), and the write to IA32\_RTIT\_CTL.TraceEn by XSAVES (Section 35.3.5.2).

### 35.2.7.2 IA32\_RTIT\_CTL MSR

IA32\_RTIT\_CTL, at address 570H, is the primary enable and control MSR for trace packet generation. Bit positions are listed in Table 35-6.

Table 35-6. IA32\_RTIT\_CTL MSR

| Position | Bit Name  | At Reset | Bit Description  |
|----------|-----------|----------|--|
| 0        | TraceEn   | 0        | If 1, enables tracing; else tracing is disabled.<br>When this bit transitions from 1 to 0, all buffered packets are flushed out of internal buffers. A further store, fence, or architecturally serializing instruction may be required to ensure that packet data can be observed at the trace endpoint. See Section 35.2.7.3 for details of enabling and disabling packet generation.<br>Note that the processor will clear this bit on #SMI (Section ) and warm reset. Other MSR bits of IA32_RTIT_CTL (and other trace configuration MSRs) are not impacted by these events. |
| 1        | CYCEn     | 0        | 0: Disables CYC Packet (see Section 35.4.2.14).<br>1: Enables CYC Packet.<br>This bit is reserved if CPUID.(EAX=14H, ECX=0):EBX.CPSB_CAM[bit 1] = 0.   |
| 2        | OS        | 0        | 0: Packet generation is disabled when CPL = 0.<br>1: Packet generation may be enabled when CPL = 0.  |
| 3        | User      | 0        | 0: Packet generation is disabled when CPL > 0.<br>1: Packet generation may be enabled when CPL > 0.  |
| 4        | PwrEvtEn  | 0        | 0: Power Event Trace packets are disabled.<br>1: Power Event Trace packets are enabled (see Section 35.2.3, “Power Event Tracing”).  |
| 5        | FUPonPTW  | 0        | 0: PTW packets are not followed by FUPs.<br>1: PTW packets are followed by FUPs.   |
| 6        | FabricEn  | 0        | 0: Trace output is directed to the memory subsystem, mechanism depends on IA32_RTIT_CTL.ToPA.<br>1: Trace output is directed to the trace transport subsystem, IA32_RTIT_CTL.ToPA is ignored.<br>This bit is reserved if CPUID.(EAX=14H, ECX=0):ECX[bit 3] = 0.  |
| 7        | CR3Filter | 0        | 0: Disables CR3 filtering.<br>1: Enables CR3 filtering.  |

Table 35-6. IA32\_RTIT\_CTL MSR (Contd.)

| Position | Bit Name  | At Reset | Bit Description  |
|----------|-----------|----------|--|
| 8        | ToPA      | 0        | 0: Single-range output scheme enabled if CPUID.(EAX=14H, ECX=0):ECX.SNGLRGNOUT[bit 2] = 1 and IA32_RTIT_CTL.FabricEn=0.<br>1: ToPA output scheme enabled (see Section 35.2.6.2) if CPUID.(EAX=14H, ECX=0):ECX.TOPA[bit 0] = 1, and IA32_RTIT_CTL.FabricEn=0.<br>Note: WRMSR to IA32_RTIT_CTL that sets TraceEn but clears this bit and FabricEn would cause #GP, if CPUID.(EAX=14H, ECX=0):ECX.SNGLRGNOUT[bit 2] = 0.<br>WRMSR to IA32_RTIT_CTL that sets this bit causes #GP, if CPUID.(EAX=14H, ECX=0):ECX.TOPA[bit 0] = 0.  |
| 9        | MTCEn     | 0        | 0: Disables MTC Packet (see Section 35.4.2.16).<br>1: Enables MTC Packet.<br>This bit is reserved if CPUID.(EAX=14H, ECX=0):EBX.MTC[bit 3] = 0.  |
| 10       | TSCEn     | 0        | 0: Disable TSC packets.<br>1: Enable TSC packets (see Section 35.4.2.11).  |
| 11       | DisRETC   | 0        | 0: Enable RET compression.<br>1: Disable RET compression (see Section 35.2.1.2).   |
| 12       | PTWEn     | 0        | 0: PTWRITE packet generation disabled.<br>1: PTWRITE packet generation enabled (see Table 35-40 “PTW Packet Definition”).  |
| 13       | BranchEn  | 0        | 0: Disable COFI-based packets.<br>1: Enable COFI-based packets: FUP, TIP, TIP.PGE, TIP.PGD, TNT, MODE.Exec, MODE.TSX.<br>See Section 35.2.5.4 for details on BranchEn.   |
| 17:14    | MTCFreq   | 0        | Defines MTC packet Frequency, which is based on the core crystal clock, or Always Running Timer (ART). MTC will be sent each time the selected ART bit toggles. The following Encodings are defined:<br>0: ART(0), 1: ART(1), 2: ART(2), 3: ART(3), 4: ART(4), 5: ART(5), 6: ART(6), 7: ART(7), 8: ART(8), 9: ART(9), 10: ART(10), 11: ART(11), 12: ART(12), 13: ART(13), 14: ART(14), 15: ART(15)<br>Software must use CPUID to query the supported encodings in the processor, see Section 35.3.1. Use of unsupported encodings will result in a #GP fault. This field is reserved if CPUID.(EAX=14H, ECX=0):EBX.MTC[bit 3] = 0.                         |
| 18       | Reserved  | 0        | Must be 0.   |
| 22:19    | CycThresh | 0        | CYC packet threshold, see Section 35.3.6 for details. CYC packets will be sent with the first eligible packet after N cycles have passed since the last CYC packet. If CycThresh is 0 then N=0, otherwise N is defined as $2^{(CycThresh-1)}$ . The following Encodings are defined:<br>0: 0, 1: 1, 2: 2, 3: 4, 4: 8, 5: 16, 6: 32, 7: 64, 8: 128, 9: 256, 10: 512, 11: 1024, 12: 2048, 13: 4096, 14: 8192, 15: 16384<br>Software must use CPUID to query the supported encodings in the processor, see Section 35.3.1. Use of unsupported encodings will result in a #GP fault. This field is reserved if CPUID.(EAX=14H, ECX=0):EBX.CPSB_CAM[bit 1] = 0. |
| 23       | Reserved  | 0        | Must be 0.   |

Table 35-6. IA32\_RTIT\_CTL MSR (Contd.)

| Position | Bit Name  | At Reset | Bit Description  |
|----------|-----------|----------|--|
| 27:24    | PSBFreq   | 0        | Indicates the frequency of PSB packets. PSB packet frequency is based on the number of Intel PT packet bytes output, so this field allows the user to determine the increment of IA32_RTIT_STATUS.PacketByteCnt that should cause a PSB to be generated. Note that PSB insertion is not precise, but the average output bytes per PSB should approximate the SW selected period. The following Encodings are defined:<br>0: 2K, 1: 4K, 2: 8K, 3: 16K, 4: 32K, 5: 64K, 6: 128K, 7: 256K, 8: 512K, 9: 1M, 10: 2M, 11: 4M, 12: 8M, 13: 16M, 14: 32M, 15: 64M<br>Software must use CPUID to query the supported encodings in the processor, see Section 35.3.1. Use of unsupported encodings will result in a #GP fault. This field is reserved if CPUID.(EAX=14H, ECX=0):EBX.CPSB_CAM[bit 1] = 0. |
| 31:28    | Reserved  | 0        | Must be 0.   |
| 35:32    | ADDR0_CFG | 0        | Configures the base/limit register pair IA32_RTIT_ADDR0_A/B based on the following encodings:<br>0: ADDR0 range unused.<br>1: The [IA32_RTIT_ADDR0_A..IA32_RTIT_ADDR0_B] range defines a FilterEn range. FilterEn will only be set when the IP is within this range, though other FilterEn ranges can additionally be used. See Section 35.2.4.3 for details on IP filtering.<br>2: The [IA32_RTIT_ADDR0_A..IA32_RTIT_ADDR0_B] range defines a TraceStop range. TraceStop will be asserted if code branches into this range. See 4.2.8 for details on TraceStop.<br>3..15: Reserved (#GP).<br>This field is reserved if CPUID.(EAX=14H, ECX=1):EBX.RANGECONT[2:0] >= 0.  |
| 39:36    | ADDR1_CFG | 0        | Configures the base/limit register pair IA32_RTIT_ADDR1_A/B based on the following encodings:<br>0: ADDR1 range unused.<br>1: The [IA32_RTIT_ADDR1_A..IA32_RTIT_ADDR1_B] range defines a FilterEn range. FilterEn will only be set when the IP is within this range, though other FilterEn ranges can additionally be used. See Section 35.2.4.3 for details on IP filtering.<br>2: The [IA32_RTIT_ADDR1_A..IA32_RTIT_ADDR1_B] range defines a TraceStop range. TraceStop will be asserted if code branches into this range. See Section 35.4.2.10 for details on TraceStop.<br>3..15: Reserved (#GP).<br>This field is reserved if CPUID.(EAX=14H, ECX=1):EBX.RANGECONT[2:0] < 2.   |
| 43:40    | ADDR2_CFG | 0        | Configures the base/limit register pair IA32_RTIT_ADDR2_A/B based on the following encodings:<br>0: ADDR2 range unused.<br>1: The [IA32_RTIT_ADDR2_A..IA32_RTIT_ADDR2_B] range defines a FilterEn range. FilterEn will only be set when the IP is within this range, though other FilterEn ranges can additionally be used. See Section 35.2.4.3 for details on IP filtering.<br>2: The [IA32_RTIT_ADDR2_A..IA32_RTIT_ADDR2_B] range defines a TraceStop range. TraceStop will be asserted if code branches into this range. See Section 35.4.2.10 for details on TraceStop.<br>3..15: Reserved (#GP).<br>This field is reserved if CPUID.(EAX=14H, ECX=1):EBX.RANGECONT[2:0] < 3.   |

Table 35-6. IA32\_RTIT\_CTL MSR (Contd.)

| Position | Bit Name             | At Reset | Bit Description   |
|----------|----------------------|----------|---|
| 47:44    | ADDR3_CFG            | 0        | Configures the base/limit register pair IA32_RTIT_ADDR3_A/B based on the following encodings:<br>0: ADDR3 range unused.<br>1: The [IA32_RTIT_ADDR3_A..IA32_RTIT_ADDR3_B] range defines a FilterEn range. FilterEn will only be set when the IP is within this range, though other FilterEn ranges can additionally be used. See Section 35.2.4.3 for details on IP filtering.<br>2: The [IA32_RTIT_ADDR3_A..IA32_RTIT_ADDR3_B] range defines a TraceStop range. TraceStop will be asserted if code branches into this range. See Section 35.4.2.10 for details on TraceStop.<br>3..15: Reserved (#GP).<br>This field is reserved if CPUID.(EAX=14H, ECX=1):EBX.RANGECNT[2:0] < 4. |
| 55:48    | Reserved             | 0        | Reserved only for future trace content enables, or address filtering configuration enables. Must be 0.  |
| 56       | InjectPsbPmiOnEnable | 0        | 1: Enables use of IA32_RTIT_STATUS bits PendPSB[6] and PendTopaPMI[7], see Section 35.2.7.4, "IA32_RTIT_STATUS MSR" for behavior of these bits.<br>0: IA32_RTIT_STATUS bits 6 and 7 are ignored.<br>This field is reserved if CPUID.(EAX=14H, ECX=0):EBX.INJECTPSBPMI[6] = 0.   |
| 59:57    | Reserved             | 0        | Reserved only for future trace content enables, or address filtering configuration enables. Must be 0.  |
| 63:60    | Reserved             | 0        | Must be 0.  |

### 35.2.7.3 Enabling and Disabling Packet Generation with TraceEn

When TraceEn transitions from 0 to 1, Intel Processor Trace is enabled, and a series of packets may be generated. These packets help ensure that the decoder is aware of the state of the processor when the trace begins, and that it can keep track of any timing or state changes that may have occurred while packet generation was disabled. A full PSB+ (see Section 35.4.2.17) will be generated if IA32\_RTIT\_STATUS.PacketByteCnt=0, and may be generated in other cases as well. Otherwise, timing packets will be generated, including TSC, TMA, and CBR (see Section 35.4.2).

In addition to the packets discussed above, if and when PacketEn (Section 35.2.5.1) transitions from 0 to 1 (which may happen immediately, depending on filtering settings), a TIP.PGE packet (Section 35.4.2.3) will be generated.

When TraceEn is set, the processor may read ToPA entries from memory and cache them internally. For this reason, software should disable packet generation before making modifications to the ToPA tables (or changing the configuration of restricted memory regions). See Section 35.7 for more details of packets that may be generated with modifications to TraceEn.

#### Disabling Packet Generation

Clearing TraceEn causes any packet data buffered within the logical processor to be flushed out, after which the output MSRs (IA32\_RTIT\_OUTPUT\_BASE and IA32\_RTIT\_OUTPUT\_MASK\_PTRS) will have stable values. When output is directed to memory, a store, fence, or architecturally serializing instruction may be required to ensure that the packet data is globally observed. No special packets are generated by disabling packet generation, though a TIP.PGD may result if PacketEn=1 at the time of disable.

#### Other Writes to IA32\_RTIT\_CTL

Any attempt to modify IA32\_RTIT\_CTL while TraceEn is set will result in a general-protection fault (#GP) unless the same write also clears TraceEn. However, writes to IA32\_RTIT\_CTL that do not modify any bits will not cause a #GP, even if TraceEn remains set.

### 35.2.7.4 IA32\_RTIT\_STATUS MSR

The IA32\_RTIT\_STATUS MSR is readable and writable by software, but some bits (ContextEn, TriggerEn) are read-only and cannot be directly modified. The WRMSR instruction ignores these bits in the source operand (attempts to modify these bits are ignored and do not cause WRMSR to fault).

This MSR can only be written when IA32\_RTIT\_CTL.TraceEn is 0; otherwise WRMSR causes a general-protection fault (#GP). The processor does not modify the value of this MSR while TraceEn is 0 (software can modify it with WRMSR).

**Table 35-7. IA32\_RTIT\_STATUS MSR**

| Position | Bit Name      | At Reset | Bit Description  |
|----------|---------------|----------|--|
| 0        | FilterEn      | 0        | This bit is written by the processor, and indicates that tracing is allowed for the current IP, see Section 35.2.5.5. Writes are ignored.  |
| 1        | ContextEn     | 0        | The processor sets this bit to indicate that tracing is allowed for the current context. See Section 35.2.5.3. Writes are ignored.   |
| 2        | TriggerEn     | 0        | The processor sets this bit to indicate that tracing is enabled. See Section 35.2.5.2. Writes are ignored.   |
| 3        | Reserved      | 0        | Must be 0.   |
| 4        | Error         | 0        | The processor sets this bit to indicate that an operational error has been encountered. When this bit is set, TriggerEn is cleared to 0 and packet generation is disabled. For details, see “ToPA Errors” in Section 35.2.6.2.<br><br>When TraceEn is cleared, software can write this bit. Once it is set, only software can clear it. It is not recommended that software ever set this bit, except in cases where it is restoring a prior saved state.  |
| 5        | Stopped       | 0        | The processor sets this bit to indicate that a ToPA Stop condition has been encountered. When this bit is set, TriggerEn is cleared to 0 and packet generation is disabled. For details, see “ToPA STOP” in Section 35.2.6.2.<br><br>When TraceEn is cleared, software can write this bit. Once it is set, only software can clear it. It is not recommended that software ever set this bit, except in cases where it is restoring a prior saved state.   |
| 6        | PendPSB       | 0        | If IA32_RTIT_CTL.InjectPsbPmiOnEnable[56] = 1, the processor sets this bit when the threshold for a PSB+ to be inserted has been reached. The processor will clear this bit when the PSB+ has been inserted into the trace. If PendPSB = 1 and InjectPsbPmiOnEnable = 1 when IA32_RTIT_CTL.TraceEn[0] transitions from 0 to 1, a PSB+ will be inserted into the trace.<br><br>This field is reserved if CPUID.(EAX=14H, ECX=0);EBX.INJECTPSBPMI[6] = 1.    |
| 7        | PendTopaPMI   | 0        | If IA32_RTIT_CTL.InjectPsbPmiOnEnable[56] = 1, the processor sets this bit when the threshold for a ToPA PMI to be inserted has been reached. Software should clear this bit once the ToPA PMI has been handled, see “ToPA PMI” for details. If PendTopaPMI = 1 and InjectPsbPmiOnEnable = 1 when IA32_RTIT_CTL.TraceEn[0] transitions from 0 to 1, a PMI will be pended.<br><br>This field is reserved if CPUID.(EAX=14H, ECX=0);EBX.INJECTPSBPMI[6] = 1. |
| 31:8     | Reserved      | 0        | Must be 0.   |
| 48:32    | PacketByteCnt | 0        | This field is written by the processor, and holds a count of packet bytes that have been sent out. The processor also uses this field to determine when the next PSB packet should be inserted. Note that the processor may clear or modify this field at any time while IA32_RTIT_CTL.TraceEn=1. It will have a stable value when IA32_RTIT_CTL.TraceEn=0. See Section 35.4.2.17 for details.   |
| 63:49    | Reserved      | 0        | Must be 0.   |

### 35.2.7.5 IA32\_RTIT\_ADDRn\_A and IA32\_RTIT\_ADDRn\_B MSRs

The role of the IA32\_RTIT\_ADDRn\_A/B register pairs, for each n, is determined by the corresponding ADDRn\_CFG fields in IA32\_RTIT\_CTL (see Section 35.2.7.2). The number of these register pairs is enumerated by CPUID.(EAX=14H, ECX=1):EAX.RANGECNT[2:0].

- Processors that enumerate support for 1 range support:  
IA32\_RTIT\_ADDR0\_A, IA32\_RTIT\_ADDR0\_B
- Processors that enumerate support for 2 ranges support:  
IA32\_RTIT\_ADDR0\_A, IA32\_RTIT\_ADDR0\_B  
IA32\_RTIT\_ADDR1\_A, IA32\_RTIT\_ADDR1\_B
- Processors that enumerate support for 3 ranges support:  
IA32\_RTIT\_ADDR0\_A, IA32\_RTIT\_ADDR0\_B  
IA32\_RTIT\_ADDR1\_A, IA32\_RTIT\_ADDR1\_B  
IA32\_RTIT\_ADDR2\_A, IA32\_RTIT\_ADDR2\_B
- Processors that enumerate support for 4 ranges support:  
IA32\_RTIT\_ADDR0\_A, IA32\_RTIT\_ADDR0\_B  
IA32\_RTIT\_ADDR1\_A, IA32\_RTIT\_ADDR1\_B  
IA32\_RTIT\_ADDR2\_A, IA32\_RTIT\_ADDR2\_B  
IA32\_RTIT\_ADDR3\_A, IA32\_RTIT\_ADDR3\_B

Each register has a single 64-bit field that holds a linear address value. Writes must ensure that the address is in canonical form, otherwise a #GP fault will result.

### 35.2.7.6 IA32\_RTIT\_CR3\_MATCH MSR

The IA32\_RTIT\_CR3\_MATCH register is compared against CR3 when IA32\_RTIT\_CTL.CR3Filter is 1. Bits 63:5 hold the CR3 address value to match, bits 4:0 are reserved to 0. For more details on CR3 filtering and the treatment of this register, see Section 35.2.4.2.

This MSR can be written only when IA32\_RTIT\_CTL.TraceEn is 0; otherwise WRMSR causes a general-protection fault (#GP). IA32\_RTIT\_CR3\_MATCH[4:0] are reserved and must be 0; an attempt to set those bits using WRMSR causes a #GP.

### 35.2.7.7 IA32\_RTIT\_OUTPUT\_BASE MSR

This MSR is used to configure the trace output destination, when output is directed to memory (IA32\_RTIT\_CTL.FabricEn = 0). The size of the address field is determined by the maximum physical address width (MAXPHYADDR), as reported by CPUID.80000008H:EAX[7:0].

When the ToPA output scheme is used, the processor may update this MSR when packet generation is enabled, and those updates are asynchronous to instruction execution. Therefore, the values in this MSR should be considered unreliable unless packet generation is disabled (IA32\_RTIT\_CTL.TraceEn = 0).

Accesses to this MSR are supported only if Intel PT output to memory is supported, hence when either CPUID.(EAX=14H, ECX=0):ECX[bit 0] or CPUID.(EAX=14H, ECX=0):ECX[bit 2] are set. Otherwise WRMSR or RDMSR cause a general-protection fault (#GP). If supported, this MSR can be written only when IA32\_RTIT\_CTL.TraceEn is 0; otherwise WRMSR causes a general-protection fault (#GP).

**Table 35-8. IA32\_RTIT\_OUTPUT\_BASE MSR**

| Position       | Bit Name     | At Reset | Bit Description   |
|----------------|--------------|----------|---|
| 6:0            | Reserved     | 0        | Must be 0.  |
| MAXPHYADDR-1:7 | BasePhysAddr | 0        | <p>The base physical address. How this address is used depends on the value of IA32_RTIT_CTL.ToPA:</p> <p>0: This is the base physical address of a single, contiguous physical output region. This could be mapped to DRAM or to MMIO, depending on the value.</p> <p>The base address should be aligned with the size of the region, such that none of the 1s in the mask value(Section 35.2.7.8) overlap with 1s in the base address. If the base is not aligned, an operational error will result (see Section 35.3.9).</p> <p>1: The base physical address of the current ToPA table. The address must be 4K aligned. Writing an address in which bits 11:7 are non-zero will not cause a #GP, but an operational error will be signaled once TraceEn is set. See “ToPA Errors” in Section 35.2.6.2 as well as Section 35.3.9.</p> |
| 63:MAXPHYADDR  | Reserved     | 0        | Must be 0.  |

**35.2.7.8 IA32\_RTIT\_OUTPUT\_MASK\_PTRS MSR**

This MSR holds any mask or pointer values needed to indicate where the next byte of trace output should be written. The meaning of the values held in this MSR depend on whether the ToPA output mechanism is in use. See Section 35.2.6.2 for details.

The processor updates this MSR while when packet generation is enabled, and those updates are asynchronous to instruction execution. Therefore, the values in this MSR should be considered unreliable unless packet generation is disabled (IA32\_RTIT\_CTL.TraceEn = 0).

Accesses to this MSR are supported only if Intel PT output to memory is supported, hence when either CPUID.(EAX=14H, ECX=0):ECX[bit 0] or CPUID.(EAX=14H, ECX=0):ECX[bit 2] are set. Otherwise WRMSR or RDMSR cause a general-protection fault (#GP). If supported, this MSR can be written only when IA32\_RTIT\_CTL.TraceEn is 0; otherwise WRMSR causes a general-protection fault (#GP).

**Table 35-9. IA32\_RTIT\_OUTPUT\_MASK\_PTRS MSR**

| Position | Bit Name          | At Reset | Bit Description   |
|----------|-------------------|----------|---|
| 6:0      | LowerMask         | 7FH      | Forced to 1, writes are ignored.  |
| 31:7     | MaskOrTableOffset | 0        | <p>The use of this field depends on the value of IA32_RTIT_CTL.ToPA:</p> <p>0: This field holds bits 31:7 of the mask value for the single, contiguous physical output region. The size of this field indicates that regions can be of size 128B up to 4GB. This value (combined with the lower 7 bits, which are reserved to 1) will be ANDed with the OutputOffset field to determine the next write address. All 1s in this field should be consecutive and starting at bit 7, otherwise the region will not be contiguous, and an operational error (Section 35.3.9) will be signaled when TraceEn is set.</p> <p>1: This field holds bits 27:3 of the offset pointer into the current ToPA table. This value can be added to the IA32_RTIT_OUTPUT_BASE value to produce a pointer to the current ToPA table entry, which itself is a pointer to the current output region. In this scenario, the lower 7 reserved bits are ignored. This field supports tables up to 256 MBytes in size.</p> |

**Table 35-9. IA32\_RTIT\_OUTPUT\_MASK\_PTRS MSR (Contd.)**

| Position | Bit Name     | At Reset | Bit Description  |
|----------|--------------|----------|--|
| 63:32    | OutputOffset | 0        | <p>The use of this field depends on the value of IA32_RTIT_CTL.ToPA:</p> <p>0: This is bits 31:0 of the offset pointer into the single, contiguous physical output region. This value will be added to the IA32_RTIT_OUTPUT_BASE value to form the physical address at which the next byte of packet output data will be written. This value must be less than or equal to the MaskOffsetTableOffset field, otherwise an operational error (Section 35.3.9) will be signaled when TraceEn is set.</p> <p>1: This field holds bits 31:0 of the offset pointer into the current ToPA output region. This value will be added to the output region base field, found in the current ToPA table entry, to form the physical address at which the next byte of trace output data will be written. This value must be less than the ToPA entry size, otherwise an operational error (Section 35.3.9) will be signaled when TraceEn is set.</p> |

## 35.2.8 Interaction of Intel® Processor Trace and Other Processor Features

### 35.2.8.1 Intel® Transactional Synchronization Extensions (Intel® TSX)

The operation of Intel TSX is described in Chapter 14 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*. For tracing purpose, packet generation does not distinguish between hardware lock elision (HLE) and restricted transactional memory (RTM), but speculative execution does have impacts on the trace output. Specifically, packets are generated as instructions complete, even for instructions in a transactional region that is later aborted. For this reason, debugging software will need indication of the beginning and end of a transactional region; this will allow software to understand when instructions are part of a transactional region and whether that region has been committed.

To enable this, TSX information is included in a MODE packet leaf. The mode bits in the leaf are:

- **InTX**: Set to 1 on an TSX transaction begin, and cleared on transaction commit or abort.
- **TXAbort**: Set to 1 only when InTX transitions from 1 to 0 on an abort. Cleared otherwise.

If BranchEn=1, this MODE packet will be sent each time the transaction status changes. See Table 35-10 for details.

**Table 35-10. TSX Packet Scenarios**

| TSX Event          | Instruction   | Packets  |
|--------------------|---|--|
| Transaction Begin  | Either XBEGIN or XACQUIRE lock (the latter if executed transactionally)   | MODE(TXAbort=0, InTX=1), FUP(CurrentIP)                |
| Transaction Commit | Either XEND or XRELEASE lock, if transactional execution ends. This happens only on the outermost commit  | MODE(TXAbort=0, InTX=0), FUP(CurrentIP)                |
| Transaction Abort  | XABORT or other transactional abort   | MODE(TXAbort=1, InTX=0), FUP(CurrentIP), TIP(TargetIP) |
| Other              | One of the following: <ul style="list-style-type: none"> <li>▪ Nested XBEGIN or XACQUIRE lock</li> <li>▪ An outer XACQUIRE lock that doesn't begin a transaction (InTX not set)</li> <li>▪ Non-outermost XEND or XRELEASE lock</li> </ul> | None. No change to TSX mode bits for these cases.      |

The CurrentIP listed above is the IP of the associated instruction. The TargetIP is the IP of the next instruction to be executed; for HLE, this is the XACQUIRE lock; for RTM, this is the fallback handler.

Intel PT stores are non-transactional, and thus packet writes are not rolled back on TSX abort.



### 35.2.8.2 TSX and IP Filtering

A complication with tracking transactions is handling transactions that start or end outside of the tracing region. Transactions can't span across a change in ContextEn, because CPL changes and CR3 changes each cause aborts. But a transaction can start within the IP filter region and end outside it.

To assist the decoder handling this situation, MODE.TSX packets can be sent even if FilterEn=0, though there will be no FUP attached. Instead, they will merely serve to indicate to the decoder when transactions are active and when they are not. When tracing resumes (due to PacketEn=1), the last MODE.TSX preceding the TIP.PGE will indicate the current transaction status.

### 35.2.8.3 System Management Mode (SMM)

SMM code has special privileges that non-SMM code does not have. Intel Processor Trace can be used to trace SMM code, but special care is taken to ensure that SMM handler context is not exposed in any non-SMM trace collection. Additionally, packet output from tracing non-SMM code cannot be written into memory space that is either protected by SMRR or used by the SMM handler.

SMM is entered via a system management interrupt (SMI). SMI delivery saves the value of IA32\_RTIT\_CTL.TraceEn into SMRAM and then clears it, thereby disabling packet generation.

The saving and clearing of IA32\_RTIT\_CTL.TraceEn ensures two things:

1. All internally buffered packet data is flushed before entering SMM (see Section 35.2.7.2).
2. Packet generation ceases before entering SMM, so any tracing that was configured outside SMM does not continue into SMM. No SMM instruction pointers or other state will be exposed in the non-SMM trace.

When the RSM instruction is executed to return from SMM, the TraceEn value that was saved by SMI delivery is restored, allowing tracing to be resumed. As is done any time packet generation is enabled, ContextEn is re-evaluated, based on the values of CPL, CR3, etc., established by RSM.

Like other interrupts, delivery of an SMI produces a FUP containing the IP of the next instruction to execute. By toggling TraceEn, SMI and RSM can produce TIP.PGD and TIP.PGE packets, respectively, indicating that tracing was disabled or re-enabled. See Table 35.7 for more information about packets entering and leaving SMM.

Although #SMI and RSM change CR3, PIP packets are not generated in these cases. With #SMI tracing is disabled before the CR3 change; with RSM TraceEn is restored after CR3 is written.

TraceEn must be cleared before executing RSM, otherwise it will cause a shutdown. Further, on processors that restrict use of Intel PT with LBRs (see Section 35.3.1.2), any RSM that results in enabling of both will cause a shutdown.

Intel PT can support tracing of System Transfer Monitor operating in SMM, see Section 35.6.

### 35.2.8.4 Virtual-Machine Extensions (VMX)

Initial implementations of Intel Processor Trace do not support tracing in VMX operation. Such processors indicate this by returning 0 for IA32\_VMX\_MISC[bit 14]. On these processors, execution of the VMXON instruction clears IA32\_RTIT\_CTL.TraceEn and any attempt to write IA32\_RTIT\_CTL in VMX operation causes a general-protection exception (#GP).

Processors that support Intel Processor Trace in VMX operation return 1 for IA32\_VMX\_MISC[bit 14]. Details of tracing in VMX operation are described in Section 35.5.

### 35.2.8.5 Intel® Software Guard Extensions (Intel® SGX)

Intel SGX provides an application with the ability to instantiate a protective container (an enclave) with confidentiality and integrity (see the *Intel® Software Guard Extensions Programming Reference*). On a processor with both Intel PT and Intel SGX enabled, when executing code within a production enclave, no control flow packets are produced by Intel PT. An enclave entry will clear ContextEn, thereby blocking control flow packet generation. A TIP.PGD packet will be generated if PacketEn=1 at the time of the entry.

Upon enclave exit, ContextEn will no longer be forced to 0. If other enables are set at the time, a TIP.PGE may be generated to indicate that tracing is resumed.

During the enclave execution, Intel PT remains enabled, and periodic or timing packets such as PSB, TSC, MTC, or CBR can still be generated. No IPs or other architectural state will be exposed.

For packet generation examples on enclave entry or exit, see Section 35.7.

### Debug Enclaves

Intel SGX allows an enclave to be configured with relaxed protection of confidentiality for debug purposes, see the *Intel® Software Guard Extensions Programming Reference*. In a debug enclave, Intel PT continues to function normally. Specifically, ContextEn is not impacted by an enclave entry or exit. Hence, the generation of ContextEn-dependent packets within a debug enclave is allowed.

#### 35.2.8.6 SENTER/ENTERACCS and ACM

GETSEC[SENDER] and GETSEC[ENTERACCS] instructions clear TraceEn, and it is not restored when those instruction complete. SENTER also causes TraceEn to be cleared on other logical processors when they rendezvous and enter the SENTER sleep state. In these two cases, the disabling of packet generation is not guaranteed to flush internally buffered packets. Some packets may be dropped.

When executing an authenticated code module (ACM), packet generation is silently disabled during ACRAM setup. TraceEn will be cleared, but no TIP.PGD packet is generated. After completion of the module, the TraceEn value will be restored. There will be no TIP.PGE packet, but timing packets, like TSC and CBR, may be produced.

#### 35.2.8.7 Intel® Memory Protection Extensions (Intel® MPX)

Bounds exceptions (#BR) caused by Intel MPX are treated like other exceptions, producing FUP and TIP packets that indicate the source and destination IPs.

## 35.3 CONFIGURATION AND PROGRAMMING GUIDELINE

### 35.3.1 Detection of Intel Processor Trace and Capability Enumeration

Processor support for Intel Processor Trace is indicated by CPUID.(EAX=07H,ECX=0H):EBX[bit 25] = 1. CPUID function 14H is dedicated to enumerate the resource and capability of processors that report CPUID.(EAX=07H,ECX=0H):EBX[bit 25] = 1. Different processor generations may have architecturally-defined variation in capabilities. Table 35-11 describes details of the enumerable capabilities that software must use across generations of processors that support Intel Processor Trace.

Table 35-11. CPUID Leaf 14H Enumeration of Intel Processor Trace Capabilities

| CPUID.(EAX=14H,ECX=0) |      | Name   | Description Behavior   |
|-----------------------|------|--|--|
| Register              | Bits |  |  |
| EAX                   | 31:0 | Maximum valid sub-leaf Index   | Specifies the index of the maximum valid sub-leaf for this CPUID leaf  |
| EBX                   | 0    | CR3 Filtering Support  | 1: Indicates that IA32_RTIT_CTL.CR3Filter can be set to 1, and that IA32_RTIT_CR3_MATCH MSR can be accessed. See Section 35.2.7.<br>0: Indicates that writes that set IA32_RTIT_CTL.CR3Filter to 1, or any access to IA32_RTIT_CR3_MATCH, will #GP fault.  |
|                       | 1    | Configurable PSB and Cycle-Accurate Mode Supported                                 | 1: (a) IA32_RTIT_CTL.PSBFreq can be set to a non-zero value, in order to select the preferred PSB frequency (see below for allowed values). (b) IA32_RTIT_STATUS.PacketByteCnt can be set to a non-zero value, and will be incremented by the processor when tracing to indicate progress towards the next PSB. If trace packet generation is enabled by setting TraceEn, a PSB will only be generated if PacketByteCnt=0. (c) IA32_RTIT_CTL.CYCEn can be set to 1 to enable Cycle-Accurate Mode. See Section 35.2.7.<br>0: (a) Any attempt to set IA32_RTIT_CTL.PSBFreq, to set IA32_RTIT_CTL.CYCEn, or write a non-zero value to IA32_RTIT_STATUS.PacketByteCnt any access to IA32_RTIT_CR3_MATCH, will #GP fault. (b) If trace packet generation is enabled by setting TraceEn, a PSB is always generated. (c) Any attempt to set IA32_RTIT_CTL.CYCEn will #GP fault.   |
|                       | 2    | IP Filtering and TraceStop supported, and Preserve Intel PT MSRs across warm reset | 1: (a) IA32_RTIT_CTL provides at one or more ADDRn_CFG field to configure the corresponding address range MSRs for IP Filtering or IP TraceStop. Each ADDRn_CFG field accepts a value in the range of 0:2 inclusive. The number of ADDRn_CFG fields is reported by CPUID.(EAX=14H, ECX=1):EAX.RANGECNT[2:0]. (b) At least one register pair IA32_RTIT_ADDRn_A and IA32_RTIT_ADDRn_B are provided to configure address ranges for IP filtering or IP TraceStop. (c) On warm reset, all Intel PT MSRs will retain their pre-reset values, though IA32_RTIT_CTL.TraceEn will be cleared. The Intel PT MSRs are listed in Section 35.2.7.<br>0: (a) An Attempt to write IA32_RTIT_CTL.ADDRn_CFG with non-zero encoding values will cause #GP. (b) Any access to IA32_RTIT_ADDRn_A and IA32_RTIT_ADDRn_B, will #GP fault. (c) On warm reset, all Intel PT MSRs will be cleared. |
|                       | 3    | MTC Supported  | 1: IA32_RTIT_CTL.MTCEn can be set to 1, and MTC packets will be generated. See Section 35.2.7.<br>0: An attempt to set IA32_RTIT_CTL.MTCEn or IA32_RTIT_CTL.MTCFreq to a non-zero value will #GP fault.  |
|                       | 4    | PTWRITE Supported  | 1: Writes can set IA32_RTIT_CTL[12] (PTWEn) and IA32_RTIT_CTL[5] (FUPonPTW), and PTWRITE can generate packets.<br>0: Writes that set IA32_RTIT_CTL[12] or IA32_RTIT_CTL[5] will #GP, and PTWRITE will #UD fault.   |
|                       | 5    | Power Event Trace Supported  | 1: Writes can set IA32_RTIT_CTL[4] (PwrEvtEn), enabling Power Event Trace packet generation.<br>0: Writes that set IA32_RTIT_CTL[4] will #GP.  |
|                       |      | 31:6   | Reserved   |

Table 35-11. CPUID Leaf 14H Enumeration of Intel Processor Trace Capabilities (Contd.)

| CPUID.(EAX=14H,ECX=0) |      | Name  | Description Behavior   |
|-----------------------|------|---|--|
| Register              | Bits |   |  |
| ECX                   | 0    | ToPA Output Supported                         | 1: Tracing can be enabled with IA32_RTIT_CTL.ToPA = 1, hence utilizing the ToPA output scheme (Section 35.2.6.2) IA32_RTIT_OUTPUT_BASE and IA32_RTIT_OUTPUT_MASK_PTRS MSRs can be accessed.<br>0: Unless CPUID.(EAX=14H, ECX=0);ECX.SNGLRNGOUT[bit 2] = 1. writes to IA32_RTIT_OUTPUT_BASE or IA32_RTIT_OUTPUT_MASK_PTRS. MSRs will #GP fault.   |
|                       | 1    | ToPA Tables Allow Multiple Output Entries     | 1: ToPA tables can hold any number of output entries, up to the maximum allowed by the MaskOffsetTableOffset field of IA32_RTIT_OUTPUT_MASK_PTRS.<br>0: ToPA tables can hold only one output entry, which must be followed by an END=1 entry which points back to the base of the table.<br>Further, ToPA PMIs will be delivered before the region is filled. See ToPA PMI in Section 35.2.6.2.<br>If there is more than one output entry before the END entry, or if the END entry has the wrong base address, an operational error will be signaled (see “ToPA Errors” in Section 35.2.6.2). |
|                       | 2    | Single-Range Output Supported                 | 1: Enabling tracing (TraceEn=1) with IA32_RTIT_CTL.ToPA=0 is supported.<br>0: Unless CPUID.(EAX=14H, ECX=0);ECX.TOPAOUT[bit 0] = 1. writes to IA32_RTIT_OUTPUT_BASE or IA32_RTIT_OUTPUT_MASK_PTRS. MSRs will #GP fault.  |
|                       | 3    | Output to Trace Transport Subsystem Supported | 1: Setting IA32_RTIT_CTL.FabricEn to 1 is supported.<br>0: IA32_RTIT_CTL.FabricEn is reserved. Write 1 to IA32_RTIT_CTL.FabricEn will #GP fault.   |
|                       | 30:4 | Reserved                                      |  |
|                       | 31   | IP Payloads are LIP                           | 1: Generated packets which contain IP payloads have LIP values, which include the CS base component.<br>0: Generated packets which contain IP payloads have RIP values, which are the offset from CS base.   |
| EDX                   | 31:0 | Reserved                                      |  |

If CPUID.(EAX=14H, ECX=0):EAX reports a non-zero value, additional capabilities of Intel Processor Trace are described in the sub-leaves of CPUID leaf 14H.

Table 35-12. CPUID Leaf 14H, sub-leaf 1H Enumeration of Intel Processor Trace Capabilities

| CPUID.(EAX=14H,ECX=1) |       | Name  | Description Behavior   |
|-----------------------|-------|---|--|
| Register              | Bits  |   |  |
| EAX                   | 2:0   | Number of Address Ranges                                | A non-zero value specifies the number ADDRn_CFG field supported in IA32_RTIT_CTL and the number of register pair IA32_RTIT_ADDRn_A/IA32_RTIT_ADDRn_B supported for IP filtering and IP TraceStop.<br><b>NOTE:</b> Currently, no processors support more than 4 address ranges.   |
|                       | 15:3  | Reserved  |  |
|                       | 31:16 | Bitmap of supported MTC Period Encodings                | The non-zero bit positions indicate the map of supported encoding values for the IA32_RTIT_CTL.MTCFreq field. This applies only if CPUID.(EAX=14H, ECX=0);EBX.MTC[bit 3] = 1 (MTC Packet generation is supported), otherwise the MTCFreq field is reserved to 0.<br>Each bit position in this field represents 1 encoding value in the 4-bit MTCFreq field (ie, bit 0 is associated with encoding value 0). For each bit:<br>1: MTCFreq can be assigned the associated encoding value.<br>0: MTCFreq cannot be assigned to the associated encoding value. A write to IA32_RTIT_CTL.MTCFreq with unsupported encoding will cause #GP fault.               |
| EBX                   | 15:0  | Bitmap of supported Cycle Threshold values              | The non-zero bit positions indicate the map of supported encoding for the IA32_RTIT_CTL.CycThresh field. This applies only if CPUID.(EAX=14H, ECX=0);EBX.CPSB_CAM[bit 1] = 1 (Cycle-Accurate Mode is Supported), otherwise the CycThresh field is reserved to 0. See Section 35.2.7.<br>Each bit position in this field represents 1 encoding value in the 4-bit CycThresh field (ie, bit 0 is associated with encoding value 0). For each bit:<br>1: CycThresh can be assigned the associated encoding value.<br>0: CycThresh cannot be assigned to the associated encoding value. A write to CycThresh with unsupported encoding will cause #GP fault. |
|                       | 31:16 | Bitmap of supported Configurable PSB Frequency encoding | The non-zero bit positions indicate the map of supported encoding for the IA32_RTIT_CTL.PSBFreq field. This applies only if CPUID.(EAX=14H, ECX=0);EBX.CPSB_CAM[bit 1] = 1 (Configurable PSB is supported), otherwise the PSBFreq field is reserved to 0. See Section 35.2.7.<br>Each bit position in this field represents 1 encoding value in the 4-bit PSBFreq field (ie, bit 0 is associated with encoding value 0). For each bit:<br>1: PSBFreq can be assigned the associated encoding value.<br>0: PSBFreq cannot be assigned to the associated encoding value. A write to PSBFreq with unsupported encoding will cause #GP fault.                |
| ECX                   | 31:0  | Reserved  |  |
| EDX                   | 31:0  | Reserved  |  |

### 35.3.1.1 Packet Decoding of RIP versus LIP

FUP, TIP, TIP.PGE, and TIP.PGE packets can contain an instruction pointer (IP) payload. On some processor generations, this payload will be an effective address (RIP), while on others this will be a linear address (LIP). In the former case, the payload is the offset from the current CS base address, while in the latter it is the sum of the offset and the CS base address (Note that in real mode, the CS base address is the value of CS<<4, while in protected mode the CS base address is the base linear address of the segment indicated by the CS register.). Which IP type is in use is indicated by enumeration (see CPUID.(EAX=14H, ECX=0):ECX.LIP[bit 31] in Table 35-11).

For software that executes while the CS base address is 0 (including all software executing in 64-bit mode), the difference is indistinguishable. A trace decoder must account for cases where the CS base address is not 0 and the resolved LIP will not be evident in a trace generated on a CPU that enumerates use of RIP. This is likely to cause problems when attempting to link the trace with the associated binaries.

Note that IP comparison logic, for IP filtering and TraceStop range calculation, is based on the same IP type as these IP packets. For processors that output RIP, the IP comparison mechanism is also based on RIP, and hence on those processors RIP values should be written to IA32\_RTIT\_ADDRn\_[AB] MSRs. This can produce differing behavior if the same trace configuration setting is run on processors reporting different IP types, i.e. CPUID.(EAX=14H, ECX=0):ECX.LIP[bit 31]. Care should be taken to check CPUID when configuring IP filters.

### 35.3.1.2 Model Specific Capability Restrictions

Some processor generations impose restrictions that prevent use of LBRs/BTS/BTM/LERs when software has enabled tracing with Intel Processor Trace. On these processors, when TraceEn is set, updates of LBR, BTS, BTM, LERs are suspended but the states of the corresponding IA32\_DEBUGCTL control fields remained unchanged as if it were still enabled. When TraceEn is cleared, the LBR array is reset, and LBR/BTS/BTM/LERs updates will resume. Further, reads of these registers will return 0, and writes will be dropped.

The list of MSRs whose updates/accesses are restricted follows.

- MSR\_LASTBRANCH\_x\_TO\_IP, MSR\_LASTBRANCH\_x\_FROM\_IP, MSR\_LBR\_INFO\_x, MSR\_LASTBRANCH\_TOS
- MSR\_LER\_FROM\_LIP, MSR\_LER\_TO\_LIP
- MSR\_LBR\_SELECT

For processor with CPUID DisplayFamily\_DisplayModel signature of 06\_3DH, 06\_47H, 06\_4EH, 06\_4FH, 06\_56H and 06\_5EH, the use of Intel PT and LBRs are mutually exclusive.

## 35.3.2 Enabling and Configuration of Trace Packet Generation

To configure trace packets, enable packet generation, and capture packets, software starts with using CPUID instruction to detect its feature flag, CPUID.(EAX=07H, ECX=0H):EBX[bit 25] = 1; followed by enumerating the capabilities described in Section 35.3.1.

Based on the capability queried from Section 35.3.1, software must configure a number of model-specific registers. This section describes programming considerations related to those MSRs.

### 35.3.2.1 Enabling Packet Generation

When configuring and enabling packet generation, the IA32\_RTIT\_CTL MSR should be written after any other Intel PT MSRs have been written, since writes to the other configuration MSRs cause a general-protection fault (#GP) if TraceEn = 1. If a prior trace collection context is not being restored, then software should first clear IA32\_RTIT\_STATUS. This is important since the Stopped, and Error fields are writable; clearing the MSR clears any values that may have persisted from prior trace packet collection contexts. See Section 35.2.7.2 for details of packets generated by setting TraceEn to 1.

If setting TraceEn to 1 causes an operational error (see Section 35.3.9), there may be a delay after the WRMSR completes before the error is signaled in the IA32\_RTIT\_STATUS MSR.

While packet generation is enabled, the values of some configuration MSRs (e.g., IA32\_RTIT\_STATUS and IA32\_RTIT\_OUTPUT\_\*) are transient, and reads may return values that are out of date. Only after packet generation is disabled (by clearing TraceEn) do reads of these MSRs return reliable values.

### 35.3.2.2 Disabling Packet Generation

After disabling packet generation by clearing IA32\_RTIT\_CTL, it is advisable to read the IA32\_RTIT\_STATUS MSR (Section 35.2.7.4):

- If the Error bit is set, an operational error was encountered, and the trace is most likely compromised. Software should check the source of the error (by examining the output MSR values), correct the source of the problem, and then attempt to gather the trace again. For details on operational errors, see Section 35.3.9. Software should clear IA32\_RTIT\_STATUS.Error before re-enabling packet generation.
- If the Stopped bit is set, software execution encountered an IP TraceStop (see Section 35.2.4.3) or the ToPA Stop condition (see “ToPA STOP” in Section 35.2.6.2) before packet generation was disabled.

### 35.3.3 Flushing Trace Output

Packets are first buffered internally and then written out asynchronously. To collect packet output for post-processing, a collector needs first to ensure that all packet data has been flushed from internal buffers. Software can ensure this by stopping packet generation by clearing IA32\_RTIT\_CTL.TraceEn (see “Disabling Packet Generation” in Section 35.2.7.2).

When software clears IA32\_RTIT\_CTL.TraceEn to flush out internally buffered packets, the logical processor issues an SFENCE operation which ensures that WC trace output stores will be ordered with respect to the next store, or serializing operation. A subsequent read from the same logical processor will see the flushed trace data, while a read from another logical processor should be preceded by a store, fence, or architecturally serializing operation on the tracing logical processor.

When the flush operations complete, the IA32\_RTIT\_OUTPUT\_\* MSR values indicate where the trace ended. While TraceEn is set, these MSRs may hold stale values. Further, if a ToPA region with INT=1 is filled, meaning a ToPA PMI has been triggered, IA32\_PERF\_GLOBAL\_STATUS.Trace\_ToPA\_PMI[55] will be set by the time the flush completes.

### 35.3.4 Warm Reset

The MSRs software uses to program Intel Processor Trace are cleared after a power-on RESET (or cold RESET). On a warm RESET, the contents of those MSRs can retain their values from before the warm RESET with the exception that IA32\_RTIT\_CTL.TraceEn will be cleared (which may have the side effect of clearing some bits in IA32\_RTIT\_STATUS).

### 35.3.5 Context Switch Consideration

To facilitate construction of instruction execution traces at the granularity of a software process or thread context, software can save and restore the states of the trace configuration MSRs across the process or thread context switch boundary. The principle is the same as saving and restoring the typical architectural processor states across context switches.

#### 35.3.5.1 Manual Trace Configuration Context Switch

The configuration can be saved and restored through a sequence of instructions of RDMSR, management of MSR content and WRMSR. To stop tracing and to ensure that all configuration MSRs contain stable values, software must clear IA32\_RTIT\_CTL.TraceEn before reading any other trace configuration MSRs. The recommended method for saving trace configuration context manually follows:

1. RDMSR IA32\_RTIT\_CTL, save value to memory
2. WRMSR IA32\_RTIT\_CTL with saved value from RDMSR above and TraceEn cleared
3. RDMSR all other configuration MSRs whose values had changed from previous saved value, save changed values to memory

When restoring the trace configuration context, IA32\_RTIT\_CTL should be restored last:

1. Read saved configuration MSR values, aside from IA32\_RTIT\_CTL, from memory, and restore them with WRMSR
2. Read saved IA32\_RTIT\_CTL value from memory, and restore with WRMSR.

### 35.3.5.2 Trace Configuration Context Switch Using XSAVES/XRSTORS

On processors whose XSAVE feature set supports XSAVES and XRSTORS, the Trace configuration state can be saved using XSAVES and restored by XRSTORS, in conjunction with the bit field associated with supervisory state component in IA32\_XSS. See Chapter 13, “Managing State Using the XSAVE Feature Set” of *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1*.

The layout of the trace configuration component state in the XSAVE area is shown in Table 35-13.<sup>1</sup>

**Table 35-13. Memory Layout of the Trace Configuration State Component**

| Offset within Component Area | Field                      | Offset within Component Area | Field                 |
|------------------------------|----------------------------|------------------------------|-----------------------|
| 0H                           | IA32_RTIT_CTL              | 08H                          | IA32_RTIT_OUTPUT_BASE |
| 10H                          | IA32_RTIT_OUTPUT_MASK_PTRS | 18H                          | IA32_RTIT_STATUS      |
| 20H                          | IA32_RTIT_CR3_MATCH        | 28H                          | IA32_RTIT_ADDR0_A     |
| 30H                          | IA32_RTIT_ADDR0_B          | 38H                          | IA32_RTIT_ADDR1_A     |
| 40H                          | IA32_RTIT_ADDR1_B          | 48H-End                      | Reserved              |

The IA32\_XSS MSR is zero coming out of RESET. Once IA32\_XSS[bit 8] is set, system software operating at CPL=0 can use XSAVES/XRSTORS with the appropriate requested-feature bitmap (RFBM) to manage supervisor state components in the XSAVE map. See Chapter 13, “Managing State Using the XSAVE Feature Set” of *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1*.

### 35.3.6 Cycle-Accurate Mode

Intel PT can be run in a cycle-accurate mode which enables CYC packets (see Section 35.4.2.14) that provide low-level information in the processor core clock domain. This cycle counter data in CYC packets can be used to compute IPC (Instructions Per Cycle), or to track wall-clock time on a fine-grain level.

To enable cycle-accurate mode packet generation, software should set IA32\_RTIT\_CTL.CYCEn=1. It is recommended that software also set TSCEn=1 anytime cycle-accurate mode is in use. With this, all CYC-eligible packets will be preceded by a CYC packet, the payload of which indicates the number of core clock cycles since the last CYC packet. In cases where multiple CYC-eligible packets are generated in a single cycle, only a single CYC will be generated before the CYC-eligible packets, otherwise each CYC-eligible packet will be preceded by its own CYC. The CYC-eligible packets are:

- TNT, TIP, TIP.PGE, TIP.PGD, MODE.EXEC, MODE.TSX, PIP, VMCS, OVF, MTC, TSC, PTWRITE, EXSTOP

TSC packets are generated when there is insufficient information to reconstruct wall-clock time, due to tracing being disabled (TriggerEn=0), or power down scenarios like a transition to a deep-sleep MWAIT C-state. In this case, the CYC that is generated along with the TSC will indicate the number of cycles actively tracing (those powered up, with TriggerEn=1) executed between the last CYC packet and the TSC packet. And hence the amount of time spent while tracing is inactive can be inferred from the difference in time between that expected based on the CYC value, and the actual time indicated by the TSC.

Additional CYC packets may be sent stand-alone, so that the processor can ensure that the decoder is aware of the number of cycles that have passed before the internal hardware counter wraps, or is reset due to other micro-architectural condition. There is no guarantee at what intervals these standalone CYC packets will be sent, except that they will be sent before the wrap occurs. An illustration is given below.

1. Table 35-13 documents support for the MSRs defining address ranges 0 and 1. Processors that provide XSAVE support for Intel Processor Trace support only those address ranges.



**Example 35-1. An Illustrative CYC Packet Example**

| Time (cycles) | Instruction Snapshot | Generated Packets | Comment   |
|---------------|----------------------|-------------------|---|
| x             | call %eax            | CYC(?), TIP       | ?Elapsed cycles from the previous CYC unknown             |
| x + 2         | call %ebx            | CYC(2), TIP       | 1 byte CYC packet; 2 cycles elapsed from the previous CYC |
| x + 8         | jnz Foo (not taken)  | CYC(6)            | 1 byte CYC packet   |
| x + 9         | ret (compressed)     |                   |   |
| x + 12        | jnz Bar (taken)      |                   |   |
| x + 16        | ret (uncompressed)   | TNT, CYC(8), TIP  | 1 byte CYC packet   |
| x + 4111      |                      | CYC(4095)         | 2 byte CYC packet   |
| x + 12305     |                      | CYC(8194)         | 3 byte CYC packet   |
| x + 16332     | mov cr3, %ebx        | CYC(4027), PIP    | 2 byte CYC packet   |

**35.3.6.1 Cycle Counter**

The cycle counter is implemented in hardware (independent of the time stamp counter or performance monitoring counters), and is a simple incrementing counter that does not saturate, but rather wraps. The size of the counter is implementation specific.

The cycle counter is reset to zero any time that TriggerEn is cleared, and when a CYC packet is sent. The cycle counter will continue to count when ContextEn or FilterEn are cleared, and cycle packets will still be generated. It will not count during sleep states that result in Intel PT logic being powered-down, but will count up to the point where clocks are disabled, and resume counting once they are re-enabled.

**35.3.6.2 Cycle Packet Semantics**

Cycle-accurate mode adheres to the following protocol:

- All packets that precede a CYC packet represent instructions or events that took place before the CYC time.
- All packets that follow a CYC packet represent instructions or events that took place at the same time as, or after, the CYC time.
- The CYC-eligible packet that immediately follows a CYC packet represents an instruction or event that took place at the same time as the CYC time.

These items above give the decoder a means to apply CYC packets to a specific instruction in the assembly stream. Most packets represent a single instruction or event, and hence the CYC packet that precedes each of those packets represents the retirement time of that instruction or event. In the case of TNT packets, up to 6 conditional branches and/or compressed RETs may be contained in the packet. In this case, the preceding CYC packet provides the retirement time of the first branch in the packet. It is possible that multiple branches retired in the same cycle as that first branch in the TNT, but the protocol will not make that obvious. Also note that a MTC packet could be generated in the same cycle as the first JCC in the TNT packet. In this case, the CYC would precede both the MTC and the TNT, and apply to both.

Note that there are times when the cycle counter will stop counting, though cycle-accurate mode is enabled. After any such scenario, a CYC packet followed by TSC packet will be sent. See Section 35.8.3.2 to understand how to interpret the payload values

**Multi-packet Instructions or Events**

Some operations, such as interrupts or task switches, generate multiple packets. In these cases, multiple CYC packets may be sent for the operation, preceding each CYC-eligible packet in the operation. An example, using a task switch on a software interrupt, is shown below.

**Example 35-2. An Example of CYC in the Presence of Multi-Packet Operations**

| Time (cycles) | Instruction Snapshot | Generated Packets                               |
|---------------|----------------------|---|
| x             | jnz Foo (not taken)  | CYC(?),   |
| x + 2         | ret (compressed)     |   |
| x + 8         | jnz Bar (taken)      |   |
| x + 9         | jmp %eax             | TNT, CYC(9), TIP                                |
| x + 12        | jnz Bar (not taken)  | CYC(3)  |
| x + 32        | int3 (task gate)     | TNT, FUP, CYC(10), PIP, CYC(20), MODE.Exec, TIP |

**35.3.6.3 Cycle Thresholds**

Software can opt to reduce the frequency of cycle packets, a trade-off to save bandwidth and intrusion at the expense of precision. This is done by utilizing a cycle threshold (see Section 35.2.7.2).

IA32\_RTIT\_CTL.CycThresh indicates to the processor the minimum number of cycles that must pass before the next CYC packet should be sent. If this value is 0, no threshold is used, and CYC packets can be sent every cycle in which a CYC-eligible packet is generated. If this value is greater than 0, the hardware will wait until the associated number of cycles have passed since the last CYC packet before sending another. CPUID provides the threshold options for CycThresh, see Section 35.3.1.

Note that the cycle threshold does not dictate how frequently a CYC packet will be posted, it merely assigns the maximum frequency. If the cycle threshold is 16, a CYC packet can be posted no more frequently than every 16 cycles. However, once that threshold of 16 cycles has passed, it still requires a new CYC-eligible packet to be generated before a CYC will be inserted. Table 35-14 illustrates the threshold behavior.

**Table 35-14. An Illustrative CYC Packet Example**

| Time (cycles) | Instruction Snapshot | Threshold |          |          |          |
|---------------|----------------------|-----------|----------|----------|----------|
|               |                      | 0         | 16       | 32       | 64       |
| x             | jmp %eax             | CYC, TIP  | CYC, TIP | CYC, TIP | CYC, TIP |
| x + 9         | call %ebx            | CYC, TIP  | TIP      | TIP      | TIP      |
| x + 15        | call %ecx            | CYC, TIP  | TIP      | TIP      | TIP      |
| x + 30        | jmp %edx             | CYC, TIP  | CYC, TIP | TIP      | TIP      |
| x + 38        | mov cr3, %eax        | CYC, PIP  | PIP      | CYC, PIP | PIP      |
| x + 46        | jmp [%eax]           | CYC, TIP  | CYC, TIP | TIP      | TIP      |
| x + 64        | call %edx            | CYC, TIP  | CYC, TIP | TIP      | CYC, TIP |
| x + 71        | jmp %edx             | CYC, TIP  | TIP      | CYC, TIP | TIP      |

**35.3.7 Decoder Synchronization (PSB+)**

The PSB packet (Section 35.4.2.17) serves as a synchronization point for a trace-packet decoder. It is a pattern in the trace log for which the decoder can quickly scan to align packet boundaries. No legal packet combination can result in such a byte sequence. As such, it serves as the starting point for packet decode. To decode a trace log properly, the decoder needs more than simply to be aligned: it needs to know some state and potentially some timing information as well. The decoder should never need to retain any information (e.g., LastIP, call stack, compound packet event) across a PSB; all compound packet events will be completed before a PSB, and any compression state will be reset.

When a PSB packet is generated, it is followed by a PSBEND packet (Section 35.4.2.18). One or more packets may be generated in between those two packets, and these inform the decoder of the current state of the processor. These packets, known collectively as PSB+, should be interpreted as “status only”, since they do not imply any change of state at the time of the PSB, nor are they associated directly with any instruction or event. Thus, the

normal binding and ordering rules that apply to these packets outside of PSB+ can be ignored when these packets are between a PSB and PSBEND. They inform the decoder of the state of the processor at the time of the PSB.

PSB+ can include:

- Timestamp (TSC), if IA32\_RTIT\_CTL.TSCEn=1.
- Timestamp-MTC Align (TMA), if IA32\_RTIT\_CTL.TSCEn=1 && IA32\_RTIT\_CTL.MTCEn=1.
- Paging Information Packet (PIP), if ContextEn=1 and IA32\_RTIT\_CTL.OS=1. The non-root bit (NR) is set if the logical processor is in VMX non-root operation and the “conceal VMX from PT” VM-execution control is 0.
- VMCS packet, if either the logical is in VMX root operation or the logical processor is in VMX non-root operation and the “conceal VMX from PT” VM-execution control is 0.
- Core Bus Ratio (CBR).
- MODE.TSX, if ContextEn=1 and BranchEn = 1.
- MODE.Exec, if PacketEn=1.
- Flow Update Packet (FUP), if PacketEn=1.

PSB is generated only when TriggerEn=1; hence PSB+ has the same dependencies. The ordering of packets within PSB+ is not fixed. Timing packets such as CYC and MTC may be generated between PSB and PSBEND, and their meanings are the same as outside PSB+.

A PSB+ can be lost in some scenarios. If IA32\_RTIT\_STATUS.TriggerEn is cleared just as the PSB threshold is reached, the PSB+ may not be generated. TriggerEn can be cleared by a WRMSR that clears IA32\_RTIT\_CTL.TraceEn, a VM-exit that clears IA32\_RTIT\_CTL.TraceEn, an #SMI, or any time that either IA32\_RTIT\_STATUS.Stopped is set (e.g., by a TraceStop or ToPA stop condition) or IA32\_RTIT\_STATUS.Error is set (e.g., by an Intel PT output error).

Note that an overflow can occur during PSB+, and this could cause the PSBEND packet to be lost. For this reason, the OVF packet should also be viewed as terminating PSB+. If IA32\_RTIT\_STATUS.TriggerEn is cleared just as the PSB threshold is reached, the PSB+ may not be generated. TriggerEn can be cleared by a WRMSR that clears IA32\_RTIT\_CTL.TraceEn, a VM-exit that clears IA32\_RTIT\_CTL.TraceEn, an #SMI, or any time that either IA32\_RTIT\_STATUS.Stopped is set (e.g., by a TraceStop or ToPA stop condition) or IA32\_RTIT\_STATUS.Error is set (e.g., by an Intel PT output error). On processors that support PSB preservation (CPUID.(EAX=14H, ECX=0):EBX.INJECTPSBPMI[6] = 1), setting IA32\_RTIT\_CTL.InjectPsbPmiOnEnable[56] = 1 will ensure that a PSB+ that is pending at the time PT is disabled will be recorded by setting IA32\_RTIT\_STATUS.PendPSB[6] = 1. A PSB will then be pended when the saved PT context is later restored.

### 35.3.8 Internal Buffer Overflow

In the rare circumstances when new packets need to be generated but the processor’s dedicated internal buffers are all full, an “internal buffer overflow” occurs. On such an overflow packet generation ceases (as packets would need to enter the processor’s internal buffer) until the overflow resolves. Once resolved, packet generation resumes.

When the buffer overflow is cleared, an OVF packet (Section 35.4.2.16) is generated, and the processor ensures that packets which follow the OVF are not compressed (IP compression or RET compression) against packets that were lost.

If IA32\_RTIT\_CTL.BranchEn = 1, the OVF packet will be followed by a FUP if the overflow resolves while PacketEn=1. If the overflow resolves while PacketEn = 0 no packet is generated, but a TIP.PGE will naturally be generated later, once PacketEn = 1. The payload of the FUP or TIP.PGE will be the Current IP of the first instruction upon which tracing resumes after the overflow is cleared. If the overflow resolves while PacketEn=1, only timing packets may come between the OVF and the FUP. If the overflow resolves while PacketEn=0, any other packets that are not dependent on PacketEn may come between the OVF and the TIP.PGE.

#### 35.3.8.1 Overflow Impact on Enables

The address comparisons to ADDRn ranges, for IP filtering and TraceStop (Section 35.2.4.3), continue during a buffer overflow, and TriggerEn, ContextEn, and FilterEn may change during a buffer overflow. Like other packets, however, any TIP.PGE or TIP.PGD packets that would have been generated will be lost. Further, IA32\_RTIT\_STATUS.PacketByteCnt will not increment, since it is only incremented when packets are generated.

If a TraceStop event occurs during the buffer overflow, IA32\_RTIT\_STATUS.Stopped will still be set, tracing will cease as a result. However, the TraceStop packet, and any TIP.PGD that result from the TraceStop, may be dropped.

### 35.3.8.2 Overflow Impact on Timing Packets

Any timing packets that are generated during a buffer overflow will be dropped. If only a few MTC packets are dropped, a decoder should be able to detect this by noticing that the time value in the first MTC packet after the buffer overflow incremented by more than one. If the buffer overflow lasted long enough that 256 MTC packets are lost (and thus the MTC packet 'wraps' its 8-bit CTC value), then the decoder may be unable to properly understand the trace. This is not an expected scenario. No CYC packets are generated during overflow, even if the cycle counter wraps.

Note that, if cycle-accurate mode is enabled, the OVF packet will generate a CYC packet. Because the cycle counter counts during overflows, this CYC packet can provide the duration of the overflow. However, there is a risk that the cycle counter wrapped during the overflow, which could render this CYC misleading.

### 35.3.9 Operational Errors

Errors are detected as a result of packet output configuration problems, which can include output alignment issues, ToPA reserved bit violations, or overlapping packet output with restricted memory. See "ToPA Errors" in Section 35.2.6.2 for details on ToPA errors, and Section 35.2.6.4 for details on restricted memory errors. Operational errors are only detected and signaled when TraceEn=1.

When an operational error is detected, tracing is disabled and the error is logged. Specifically, IA32\_RTIT\_STATUS.Error is set, which will cause IA32\_RTIT\_STATUS.TriggerEn to be 0. This will disable generation of all packets. Some causes of operational errors may lead to packet bytes being dropped.

It should be noted that the timing of error detection may not be predictable. Errors are signaled when the processor encounters the problematic configuration. This could be as soon as packet generation is enabled but could also be later when the problematic entry or field needs to be used.

Once an error is signaled, software should disable packet generation by clearing TraceEn, diagnose and fix the error condition, and clear IA32\_RTIT\_STATUS.Error. At this point, packet generation can be re-enabled.

## 35.4 TRACE PACKETS AND DATA TYPES

This section details the data packets generated by Intel Processor Trace. It is useful for developers writing the interpretation code that will decode the data packets and apply it to the traced source code.

### 35.4.1 Packet Relationships and Ordering

This section introduces the concept of packet "binding", which involves determining the IP in a binary disassembly at which the change indicated by a given packet applies. Some packets have the associated IP as the payload (FUP, TIP), while for others the decoder need only search for the next instance of a particular instruction (or instructions) to bind the packet (TNT). However, in many cases, the decoder will need to consider the relationship between packets, and to use this packet context to determine how to bind the packet.

Section 35.4.2 below provides detailed descriptions of the packets, including how packets bind to IPs in the disassembly, to other packets, or to nothing at all. Many packets listed are simple to bind, because they are generated in only a few scenarios. Those that require more consideration are typically part of "compound packet events", such as interrupts, exceptions, and some instructions, where multiple packets are generated by a single operation (instruction or event). These compound packet events frequently begin with a FUP to indicate the source address (if it is not clear from the disassembly), and are concluded by a TIP or TIP.PGD packet that indicates the destination address (if one is provided). In this scenario, the FUP is said to be "coupled" with the TIP packet.

Other packets could be in between the coupled FUP and TIP packet. Timing packets, such as TSC, MTC, CYC, or CBR, could arrive at any time, and hence could intercede in a compound packet event. If an operation changes CR3 or the processor's mode of execution, a state update packet (i.e., PIP or MODE) is generated. The state changes

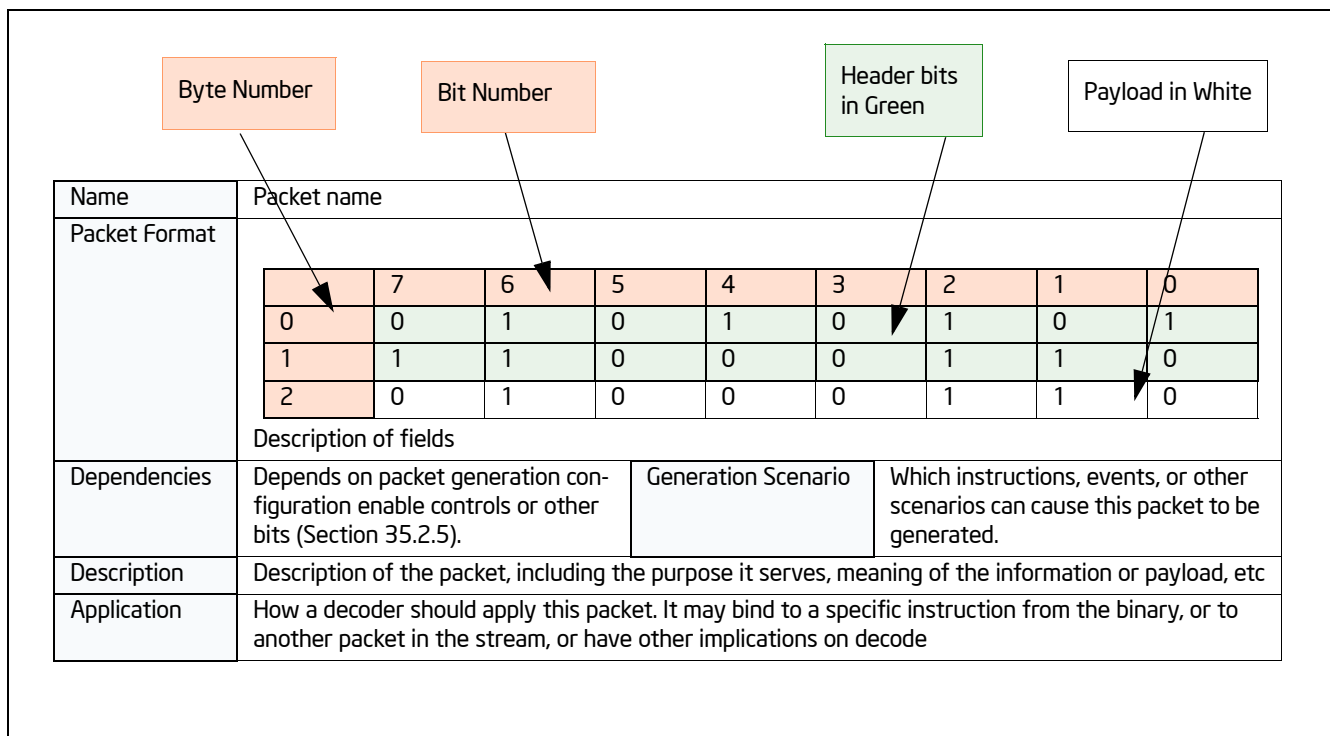
indicated by these intermediate packets should be applied at the IP of the TIP\* packet. A summary of compound packet events is provided in Table 35-15; see Section 35.4.2 for more per-packet details and Section 35.7 for more detailed packet generation examples.

**Table 35-15. Compound Packet Event Summary**

| Event Type  | Beginning                   | Middle   | End                   | Comment  |
|---|-----------------------------|--|-----------------------|--|
| Unconditional, uncompressed control-flow transfer | FUP or none                 | Any combination of PIP, VMCS, MODE.Exec, or none | TIP or TIP.PGD        | FUP only for asynchronous events. Order of middle packets may vary.<br>PIP/VMCS/MODE only if the operation modifies the state tracked by these respective packets. |
| TSX Update  | MODE.TSX, and (FUP or none) | None   | TIP, TIP.PGD, or none | FUP<br>TIP/TIP.PGD only for TSX abort cases.   |
| Overflow  | OVF                         | PSB, PSBEND, or none                             | FUP or TIP.PGE        | FUP if overflow resolves while ContextEn=1, else TIP.PGE.  |

### 35.4.2 Packet Definitions

The following description of packet definitions are in tabular format. Figure 35-3 explains how to interpret them. Packet bits listed as “RSVD” are not guaranteed to be 0.



**Figure 35-3. Interpreting Tabular Definition of Packet Format**

### 35.4.2.1 Taken/Not-taken (TNT) Packet

Table 35-16. TNT Packet Definition

|               |  |                 |                 |                     |                 |   |                 |                 |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|---------------|--|-----------------|-----------------|---------------------|-----------------|---|-----------------|-----------------|-----------|--|--|---|---|---|---|---|---|---|---|--|---|---|----------------|----------------|----------------|----------------|----------------|----------------|---|-----------|---|---|---|---|---|---|---|---|---|---|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|---|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|---|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|---|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|---|----------------|----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|---|---|----------------|----------------|----------------|----------------|----------------|----------------|----------------|
| Name          | Taken/Not-taken (TNT) Packet   |                 |                 |                     |                 |   |                 |                 |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| Packet Format | <table border="1" style="width:100%; text-align:center;"> <tr> <td></td><td>7</td><td>6</td><td>5</td><td>4</td><td>3</td><td>2</td><td>1</td><td>0</td><td></td></tr> <tr> <td>0</td><td>1</td><td>B<sub>1</sub></td><td>B<sub>2</sub></td><td>B<sub>3</sub></td><td>B<sub>4</sub></td><td>B<sub>5</sub></td><td>B<sub>6</sub></td><td>0</td><td>Short TNT</td></tr> </table>   |                 |                 |                     |                 |   |                 |                 |           |  |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |  | 0 | 1 | B <sub>1</sub> | B <sub>2</sub> | B <sub>3</sub> | B <sub>4</sub> | B <sub>5</sub> | B <sub>6</sub> | 0 | Short TNT |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               |  | 7               | 6               | 5                   | 4               | 3   | 2               | 1               | 0         |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 0             | 1  | B <sub>1</sub>  | B <sub>2</sub>  | B <sub>3</sub>      | B <sub>4</sub>  | B <sub>5</sub>  | B <sub>6</sub>  | 0               | Short TNT |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               | B1...BN represent the last N conditional branch or compressed RET (Section 35.4.2.2) results, such that B1 is oldest and BN is youngest. The short TNT packet can contain from 1 to 6 TNT bits. The long TNT packet can contain from 1 to 47 TNT bits.   |                 |                 |                     |                 |   |                 |                 |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               | <table border="1" style="width:100%; text-align:center;"> <tr> <td></td><td>7</td><td>6</td><td>5</td><td>4</td><td>3</td><td>2</td><td>1</td><td>0</td><td></td></tr> <tr> <td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td rowspan="8">Long TNT</td></tr> <tr> <td>1</td><td>1</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>1</td></tr> <tr> <td>2</td><td>B<sub>40</sub></td><td>B<sub>41</sub></td><td>B<sub>42</sub></td><td>B<sub>43</sub></td><td>B<sub>44</sub></td><td>B<sub>45</sub></td><td>B<sub>46</sub></td><td>B<sub>47</sub></td></tr> <tr> <td>3</td><td>B<sub>32</sub></td><td>B<sub>33</sub></td><td>B<sub>34</sub></td><td>B<sub>35</sub></td><td>B<sub>36</sub></td><td>B<sub>37</sub></td><td>B<sub>38</sub></td><td>B<sub>39</sub></td></tr> <tr> <td>4</td><td>B<sub>24</sub></td><td>B<sub>25</sub></td><td>B<sub>26</sub></td><td>B<sub>27</sub></td><td>B<sub>28</sub></td><td>B<sub>29</sub></td><td>B<sub>30</sub></td><td>B<sub>31</sub></td></tr> <tr> <td>5</td><td>B<sub>16</sub></td><td>B<sub>17</sub></td><td>B<sub>18</sub></td><td>B<sub>19</sub></td><td>B<sub>20</sub></td><td>B<sub>21</sub></td><td>B<sub>22</sub></td><td>B<sub>23</sub></td></tr> <tr> <td>6</td><td>B<sub>8</sub></td><td>B<sub>9</sub></td><td>B<sub>10</sub></td><td>B<sub>11</sub></td><td>B<sub>12</sub></td><td>B<sub>13</sub></td><td>B<sub>14</sub></td><td>B<sub>15</sub></td></tr> <tr> <td>7</td><td>1</td><td>B<sub>1</sub></td><td>B<sub>2</sub></td><td>B<sub>3</sub></td><td>B<sub>4</sub></td><td>B<sub>5</sub></td><td>B<sub>6</sub></td><td>B<sub>7</sub></td></tr> </table> |                 |                 |                     |                 |   |                 |                 |           |  |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |  | 0 | 0 | 0              | 0              | 0              | 0              | 0              | 1              | 0 | Long TNT  | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 1 | 2 | B <sub>40</sub> | B <sub>41</sub> | B <sub>42</sub> | B <sub>43</sub> | B <sub>44</sub> | B <sub>45</sub> | B <sub>46</sub> | B <sub>47</sub> | 3 | B <sub>32</sub> | B <sub>33</sub> | B <sub>34</sub> | B <sub>35</sub> | B <sub>36</sub> | B <sub>37</sub> | B <sub>38</sub> | B <sub>39</sub> | 4 | B <sub>24</sub> | B <sub>25</sub> | B <sub>26</sub> | B <sub>27</sub> | B <sub>28</sub> | B <sub>29</sub> | B <sub>30</sub> | B <sub>31</sub> | 5 | B <sub>16</sub> | B <sub>17</sub> | B <sub>18</sub> | B <sub>19</sub> | B <sub>20</sub> | B <sub>21</sub> | B <sub>22</sub> | B <sub>23</sub> | 6 | B <sub>8</sub> | B <sub>9</sub> | B <sub>10</sub> | B <sub>11</sub> | B <sub>12</sub> | B <sub>13</sub> | B <sub>14</sub> | B <sub>15</sub> | 7 | 1 | B <sub>1</sub> | B <sub>2</sub> | B <sub>3</sub> | B <sub>4</sub> | B <sub>5</sub> | B <sub>6</sub> | B <sub>7</sub> |
|               | 7  | 6               | 5               | 4                   | 3               | 2   | 1               | 0               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 0             | 0  | 0               | 0               | 0                   | 0               | 0   | 1               | 0               | Long TNT  |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 1             | 1  | 0               | 1               | 0                   | 0               | 0   | 1               | 1               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 2             | B <sub>40</sub>  | B <sub>41</sub> | B <sub>42</sub> | B <sub>43</sub>     | B <sub>44</sub> | B <sub>45</sub>   | B <sub>46</sub> | B <sub>47</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 3             | B <sub>32</sub>  | B <sub>33</sub> | B <sub>34</sub> | B <sub>35</sub>     | B <sub>36</sub> | B <sub>37</sub>   | B <sub>38</sub> | B <sub>39</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 4             | B <sub>24</sub>  | B <sub>25</sub> | B <sub>26</sub> | B <sub>27</sub>     | B <sub>28</sub> | B <sub>29</sub>   | B <sub>30</sub> | B <sub>31</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 5             | B <sub>16</sub>  | B <sub>17</sub> | B <sub>18</sub> | B <sub>19</sub>     | B <sub>20</sub> | B <sub>21</sub>   | B <sub>22</sub> | B <sub>23</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 6             | B <sub>8</sub>   | B <sub>9</sub>  | B <sub>10</sub> | B <sub>11</sub>     | B <sub>12</sub> | B <sub>13</sub>   | B <sub>14</sub> | B <sub>15</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 7             | 1  | B <sub>1</sub>  | B <sub>2</sub>  | B <sub>3</sub>      | B <sub>4</sub>  | B <sub>5</sub>  | B <sub>6</sub>  | B <sub>7</sub>  |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               | <p>Irrespective of how many TNT bits is in a packet, the last valid TNT bit is followed by a trailing 1, or Stop bit, as shown above. If the TNT packet is not full (fewer than 6 TNT bits for the Short TNT, or fewer than 47 TNT bits for the Long TNT), the Stop bit moves up, and the trailing bits of the packet are filled with 0s. Examples of these "partial TNTs" are shown below.</p> <table border="1" style="width:100%; text-align:center;"> <tr> <td></td><td>7</td><td>6</td><td>5</td><td>4</td><td>3</td><td>2</td><td>1</td><td>0</td><td></td></tr> <tr> <td>0</td><td>0</td><td>0</td><td>1</td><td>B<sub>1</sub></td><td>B<sub>2</sub></td><td>B<sub>3</sub></td><td>B<sub>4</sub></td><td>0</td><td>Short TNT</td></tr> </table>   |                 |                 |                     |                 |   |                 |                 |           |  |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |  | 0 | 0 | 0              | 1              | B <sub>1</sub> | B <sub>2</sub> | B <sub>3</sub> | B <sub>4</sub> | 0 | Short TNT |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               | 7  | 6               | 5               | 4                   | 3               | 2   | 1               | 0               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 0             | 0  | 0               | 1               | B <sub>1</sub>      | B <sub>2</sub>  | B <sub>3</sub>  | B <sub>4</sub>  | 0               | Short TNT |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
|               | <table border="1" style="width:100%; text-align:center;"> <tr> <td></td><td>7</td><td>6</td><td>5</td><td>4</td><td>3</td><td>2</td><td>1</td><td>0</td><td></td></tr> <tr> <td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td rowspan="8">Long TNT</td></tr> <tr> <td>1</td><td>1</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>1</td></tr> <tr> <td>2</td><td>B<sub>24</sub></td><td>B<sub>25</sub></td><td>B<sub>26</sub></td><td>B<sub>27</sub></td><td>B<sub>28</sub></td><td>B<sub>29</sub></td><td>B<sub>30</sub></td><td>B<sub>31</sub></td></tr> <tr> <td>3</td><td>B<sub>16</sub></td><td>B<sub>17</sub></td><td>B<sub>18</sub></td><td>B<sub>19</sub></td><td>B<sub>20</sub></td><td>B<sub>21</sub></td><td>B<sub>22</sub></td><td>B<sub>23</sub></td></tr> <tr> <td>4</td><td>B<sub>8</sub></td><td>B<sub>9</sub></td><td>B<sub>10</sub></td><td>B<sub>11</sub></td><td>B<sub>12</sub></td><td>B<sub>13</sub></td><td>B<sub>14</sub></td><td>B<sub>15</sub></td></tr> <tr> <td>5</td><td>1</td><td>B<sub>1</sub></td><td>B<sub>2</sub></td><td>B<sub>3</sub></td><td>B<sub>4</sub></td><td>B<sub>5</sub></td><td>B<sub>6</sub></td><td>B<sub>7</sub></td></tr> <tr> <td>6</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr> <td>7</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr> </table>   |                 |                 |                     |                 |   |                 |                 |           |  |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |  | 0 | 0 | 0              | 0              | 0              | 0              | 0              | 1              | 0 | Long TNT  | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 1 | 2 | B <sub>24</sub> | B <sub>25</sub> | B <sub>26</sub> | B <sub>27</sub> | B <sub>28</sub> | B <sub>29</sub> | B <sub>30</sub> | B <sub>31</sub> | 3 | B <sub>16</sub> | B <sub>17</sub> | B <sub>18</sub> | B <sub>19</sub> | B <sub>20</sub> | B <sub>21</sub> | B <sub>22</sub> | B <sub>23</sub> | 4 | B <sub>8</sub>  | B <sub>9</sub>  | B <sub>10</sub> | B <sub>11</sub> | B <sub>12</sub> | B <sub>13</sub> | B <sub>14</sub> | B <sub>15</sub> | 5 | 1               | B <sub>1</sub>  | B <sub>2</sub>  | B <sub>3</sub>  | B <sub>4</sub>  | B <sub>5</sub>  | B <sub>6</sub>  | B <sub>7</sub>  | 6 | 0              | 0              | 0               | 0               | 0               | 0               | 0               | 0               | 7 | 0 | 0              | 0              | 0              | 0              | 0              | 0              | 0              |
|               | 7  | 6               | 5               | 4                   | 3               | 2   | 1               | 0               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 0             | 0  | 0               | 0               | 0                   | 0               | 0   | 1               | 0               | Long TNT  |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 1             | 1  | 0               | 1               | 0                   | 0               | 0   | 1               | 1               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 2             | B <sub>24</sub>  | B <sub>25</sub> | B <sub>26</sub> | B <sub>27</sub>     | B <sub>28</sub> | B <sub>29</sub>   | B <sub>30</sub> | B <sub>31</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 3             | B <sub>16</sub>  | B <sub>17</sub> | B <sub>18</sub> | B <sub>19</sub>     | B <sub>20</sub> | B <sub>21</sub>   | B <sub>22</sub> | B <sub>23</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 4             | B <sub>8</sub>   | B <sub>9</sub>  | B <sub>10</sub> | B <sub>11</sub>     | B <sub>12</sub> | B <sub>13</sub>   | B <sub>14</sub> | B <sub>15</sub> |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 5             | 1  | B <sub>1</sub>  | B <sub>2</sub>  | B <sub>3</sub>      | B <sub>4</sub>  | B <sub>5</sub>  | B <sub>6</sub>  | B <sub>7</sub>  |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 6             | 0  | 0               | 0               | 0                   | 0               | 0   | 0               | 0               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| 7             | 0  | 0               | 0               | 0                   | 0               | 0   | 0               | 0               |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |
| Dependencies  | PacketEn   |                 |                 | Generation Scenario |                 | On a conditional branch or compressed RET, if it fills the TNT. Also, partial TNTs may be generated at any time, as a result of other packets being generated, or certain micro-architectural conditions occurring, before the TNT is full. |                 |                 |           |  |  |   |   |   |   |   |   |   |   |  |   |   |                |                |                |                |                |                |   |           |   |   |   |   |   |   |   |   |   |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                 |                 |                 |                 |                 |                 |                 |                 |   |                |                |                 |                 |                 |                 |                 |                 |   |   |                |                |                |                |                |                |                |

**Table 35-16. TNT Packet Definition (Contd.)**

|             |  |
|-------------|--|
| Description | Provides the taken/not-taken results for the last 1-N conditional branches (Jcc, J*CXZ, or LOOP) or compressed RETs (Section 35.4.2.2). The TNT payload bits should be interpreted as follows: <ul style="list-style-type: none"> <li>▪ 1 indicates a taken conditional branch, or a compressed RET</li> <li>▪ 0 indicates a not-taken conditional branch</li> </ul> |
| Application | Each valid payload bit (that is, bits between the header bits and the trailing Stop bit) applies to an upcoming conditional branch or RET instruction. Once a decoder consumes a TNT packet with N valid payload bits, these bits should be applied to (and hence provide the destination for) the next N conditional branches or RETs                               |

### 35.4.2.2 Target IP (TIP) Packet

**Table 35-17. IP Packet Definition**

| Name          | Target IP (TIP) Packet  |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
|---------------|---|---------------------|---|---|---|---|---|---|---|---|---|---|---|---------|--|--|---|---|---|---|---|---|---------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%;"></th> <th style="width: 15%;">7</th> <th style="width: 15%;">6</th> <th style="width: 15%;">5</th> <th style="width: 15%;">4</th> <th style="width: 15%;">3</th> <th style="width: 15%;">2</th> <th style="width: 15%;">1</th> <th style="width: 15%;">0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td colspan="3">IPBytes</td> <td>0</td> <td>1</td> <td>1</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="8">TargetIP[7:0]</td> </tr> <tr> <td>2</td> <td colspan="8">TargetIP[15:8]</td> </tr> <tr> <td>3</td> <td colspan="8">TargetIP[23:16]</td> </tr> <tr> <td>4</td> <td colspan="8">TargetIP[31:24]</td> </tr> <tr> <td>5</td> <td colspan="8">TargetIP[39:32]</td> </tr> <tr> <td>6</td> <td colspan="8">TargetIP[47:40]</td> </tr> <tr> <td>7</td> <td colspan="8">TargetIP[55:48]</td> </tr> <tr> <td>8</td> <td colspan="8">TargetIP[63:56]</td> </tr> </tbody> </table> |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | IPBytes |  |  | 0 | 1 | 1 | 0 | 1 | 1 | TargetIP[7:0] |  |  |  |  |  |  |  | 2 | TargetIP[15:8] |  |  |  |  |  |  |  | 3 | TargetIP[23:16] |  |  |  |  |  |  |  | 4 | TargetIP[31:24] |  |  |  |  |  |  |  | 5 | TargetIP[39:32] |  |  |  |  |  |  |  | 6 | TargetIP[47:40] |  |  |  |  |  |  |  | 7 | TargetIP[55:48] |  |  |  |  |  |  |  | 8 | TargetIP[63:56] |  |  |  |  |  |  |  |
|               | 7   | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 0             | IPBytes   |                     |   | 0 | 1 | 1 | 0 | 1 |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 1             | TargetIP[7:0]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 2             | TargetIP[15:8]  |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 3             | TargetIP[23:16]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 4             | TargetIP[31:24]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 5             | TargetIP[39:32]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 6             | TargetIP[47:40]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 7             | TargetIP[55:48]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 8             | TargetIP[63:56]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Dependencies  | PacketEn  | Generation Scenario | Indirect branch (including un-compressed RET), far branch, interrupt, exception, INIT, SIPI, VM exit, VM entry, TSX abort, EENTER, EEXIT, ERESUME, AEX <sup>1</sup> . |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Description   | Provides the target for some control flow transfers   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Application   | <p>Anytime a TIP is encountered, it indicates that control was transferred to the IP provided in the payload.</p> <p>The source of this control flow change, and hence the IP or instruction to which it binds, depends on the packets that precede the TIP. If a TIP is encountered and all preceding packets have already been bound, then the TIP will apply to the upcoming indirect branch, far branch, or VMRESUME. However, if there was a preceding FUP that remains unbound, it will bind to the TIP. Here, the TIP provides the target of an asynchronous event or TSX abort that occurred at the IP given in the FUP payload. Note that there may be other packets, in addition to the FUP, which will bind to the TIP packet. See the packet application descriptions for other packets for details.</p>  |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |

**NOTES:**

1. EENTER, EEXIT, ERESUME, AEX would be possible only for a debug enclave.

### IP Compression

The IP payload in a TIP, FUP, TIP.PGE, or TIP.PGD packet can vary in size, based on the mode of execution, and the use of IP compression. IP compression is an optional compression technique the processor may choose to employ to reduce bandwidth. With IP compression, the IP to be represented in the payload is compared with the last IP sent out, via any of FUP, TIP, TIP.PGE, or TIP.PGD. If that previous IP had the same upper (most significant) address bytes, those matching bytes may be suppressed in the current packet. The processor maintains an internal state of the "Last IP" that was encoded in trace packets, thus the decoder will need to keep track of the "Last IP" state in software, to match fidelity with packets generated by hardware. "Last IP" is initialized to zero, hence if the first IP in the trace may be compressed if the upper bytes are zeroes.

The “IPBytes” field of the IP packets (FUP, TIP, TIP.PGE, TIP.PGD) serves to indicate how many bytes of payload are provided, and how the decoder should fill in any suppressed bytes. The algorithm for reconstructing the IP for a TIP/FUP packet is shown in the table below.

**Table 35-18. FUP/TIP IP Reconstruction**

| IPBytes | Uncompressed IP Value      |       |                  |       |                  |       |                  |     |
|---------|----------------------------|-------|------------------|-------|------------------|-------|------------------|-----|
|         | 63:56                      | 55:48 | 47:40            | 39:32 | 31:24            | 23:16 | 15:8             | 7:0 |
| 000b    | None, IP is out of context |       |                  |       |                  |       |                  |     |
| 001b    | Last IP[63:16]             |       |                  |       |                  |       | IP Payload[15:0] |     |
| 010b    | Last IP[63:32]             |       |                  |       | IP Payload[31:0] |       |                  |     |
| 011b    | IP Payload[47] extended    |       | IP Payload[47:0] |       |                  |       |                  |     |
| 100b    | Last IP [63:48]            |       | IP Payload[47:0] |       |                  |       |                  |     |
| 101b    | Reserved                   |       |                  |       |                  |       |                  |     |
| 110b    | IP Payload[63:0]           |       |                  |       |                  |       |                  |     |
| 111b    | Reserved                   |       |                  |       |                  |       |                  |     |

The processor-internal Last IP state is guaranteed to be reset to zero when a PSB is sent out. This means that the IP that follows the PSB with either be un-compressed (011b or 110b, see Table 35-18), or compressed against zero.

At times, “IPbytes” will have a value of 0. As shown above, this does not mean that the IP payload matches the full address of the last IP, but rather that the IP for this packet was suppressed. This is used for cases where the IP that applies to the packet is out of context. An example is the TIP.PGD sent on a SYSCALL, when tracing only USR code. In that case, no TargetIP will be included in the packet, since that would expose an instruction point at CPL = 0. When the IP payload is suppressed in this manner, Last IP is not cleared, and instead refers to the last IP packet with a non-zero IPBytes field.

On processors that support a maximum linear address size of 32 bits, IP payloads may never exceed 32 bits (IPBytes <= 010b).

### Indirect Transfer Compression for Returns (RET)

In addition to IP compression, TIP packets for near return (RET) instructions can also be compressed. If the RET target matches the next IP of the corresponding CALL, then the TIP packet is unneeded, since the decoder can deduce the target IP by maintaining a CALL/RET stack of its own.

A CALL/RET stack can be maintained by the decoder by doing the following:

1. Allocate space to store 64 RET targets.
2. For near CALLs, push the Next IP onto the stack. Once the stack is full, new CALLs will force the oldest entry off the end of the stack, such that only the youngest 64 entries are stored. Note that this excludes zero-length CALLs, which are direct near CALLs with displacement zero (to the next IP). These CALLs typically don't have matching RETs.
3. For near RETs, pop the top (youngest) entry off the stack. This will be the target of the RET.

In cases where the RET is compressed, the target is guaranteed to match the value produced in 2) above. If the target is not compressed, a TIP packet will be generated with the RET target, which may differ from 2).

The hardware ensure that packets read by the decoder will always have seen the CALL that corresponds to any compressed RET. The processor will never compress a RET across a PSB, a buffer overflow, or scenario where PacketEn=0. This means that a RET whose corresponding CALL executed while PacketEn=0, or before the last PSB, etc., will not be compressed.

If the CALL/RET stack is manipulated or corrupted by software, and thereby causes a RET to transfer control to a target that is inconsistent with the CALL/RET stack, then the RET will not be compressed, and will produce a TIP packet. This can happen, for example, if software executes a PUSH instruction to push a target onto the stack, and a later RET uses this target.



When a RET is compressed, a Taken indication is added to the TNT buffer. Because it sends no TIP packet, it also does not update the internal Last IP value, and thus the decoder should treat it the same way. If the RET is not compressed, it will generate a TIP packet (just like when RET compression is disabled, via IA32\_RTIT\_CTL.DisRETC). For processors that employ deferred TIPs (Section 35.4.2.3), an uncompressed RET will not be deferred, and hence will force out any accumulated TNTs or TIPs. This serves to avoid ambiguity, and make clear to the decoder whether the near RET was compressed, and hence a bit in the in-progress TNT should be consumed, or uncompressed, in which case there will be no in-progress TNT and thus a TIP should be consumed.

Note that in the unlikely case that a RET executes in a different execution mode than the associated CALL, the decoder will need to model the same behavior with its CALL stack. For instance, if a CALL executes in 64-bit mode, a 64-bit IP value will be pushed onto the software stack. If the corresponding RET executes in 32-bit mode, then only the lower 32 target bits will be popped off of the stack, which may mean that the RET does not go to the CALL's Next IP. This is architecturally correct behavior, and this RET could be compressed, thus the decoder should match this behavior

### 35.4.2.3 Deferred TIPs

The processor may opt to defer sending out the TNT when TIPs are generated. Thus, rather than sending a partial TNT followed by a TIP, both packets will be deferred while the TNT accumulates more Jcc/RET results. Any number of TIP packets may be accumulated this way, such that only once the TNT is filled, or once another packet (e.g., FUP) is generated, the TNT will be sent, followed by all the deferred TIP packets, and finally terminated by the other packet(s) that forced out the TNT and TIP packets. Generation of many other packets (see list below) will force out the TNT and any accumulated TIP packets. This is an optional optimization in hardware to reduce the bandwidth consumption, and hence the performance impact, incurred by tracing.

**Table 35-19. TNT Examples with Deferred TIPs**

| Code Flow  | Packets, Non-Deferred TIPs            | Packets, Deferred TIPs   |
|--|---------------------------------------|--|
| 0x1000 cmp %rcx, 0<br>0x1004 jnz Foo // not-taken<br>0x1008 jmp %rdx   | TNT(0b0), TIP(0x1308)                 |  |
| 0x1308 cmp %rcx, 1<br>0x130c jnz Bar // not-taken<br>0x1310 cmp %rcx, 2<br>0x1314 jnz Baz // taken<br>0x1500 cmp %eax, 7<br>0x1504 jg Exit // not-taken<br>0x1508 jmp %r15 | TNT(0b010), TIP(0x1100)               |  |
| 0x1100 cmp %rbx, 1<br>0x1104 jg Start // not-taken<br>0x1108 add %rcx, %eax<br>0x110c ... // <b>an asynchronous interrupt arrives</b><br>INThandler:<br>0xcc00 pop %rdx    | TNT(0b0), FUP(0x110c),<br>TIP(0xcc00) | TNT(0b00100), TIP(0x1308),<br>TIP(0x1100), FUP(0x110c),<br>TIP(0xcc00) |

### 35.4.2.4 Packet Generation Enable (TIP.PGE) Packet

Table 35-20. TIP.PGE Packet Definition

|               |  |                 |   |                     |  |   |   |   |   |
|---------------|--|-----------------|---|---------------------|--|---|---|---|---|
| Name          | Target IP - Packet Generation Enable (TIP.PGE) Packet  |                 |   |                     |  |   |   |   |   |
| Packet Format |  | 7               | 6 | 5                   | 4  | 3 | 2 | 1 | 0 |
|               | 0  | IPBytes         |   |                     | 1  | 0 | 0 | 0 | 1 |
|               | 1  | TargetIP[7:0]   |   |                     |  |   |   |   |   |
|               | 2  | TargetIP[15:8]  |   |                     |  |   |   |   |   |
|               | 3  | TargetIP[23:16] |   |                     |  |   |   |   |   |
|               | 4  | TargetIP[31:24] |   |                     |  |   |   |   |   |
|               | 5  | TargetIP[39:32] |   |                     |  |   |   |   |   |
|               | 6  | TargetIP[47:40] |   |                     |  |   |   |   |   |
|               | 7  | TargetIP[55:48] |   |                     |  |   |   |   |   |
|               | 8  | TargetIP[63:56] |   |                     |  |   |   |   |   |
| Dependencies  | PacketEn transitions to 1  |                 |   | Generation Scenario | Any branch instruction, control flow transfer, or MOV CR3 that sets PacketEn, a WRMSR that enables packet generation and sets PacketEn |   |   |   |   |
| Description   | <p>Indicates that PacketEn has transitioned to 1. It provides the IP at which the tracing begins. This can occur due to any of the enables that comprise PacketEn transitioning from 0 to 1, as long as all the others are asserted. Examples:</p> <ul style="list-style-type: none"> <li>▪ TriggerEn: This is set on software write to set IA32_RTIT_CTL.TraceEn as long as the Stopped and Error bits in IA32_RTIT_STATUS are clear. The IP payload will be the Next IP of the WRMSR.</li> <li>▪ FilterEn: This is set when software jumps into the tracing region. This region is defined by enabling IP filtering in IA32_RTIT_CTL.ADDRn_CFG, and defining the range in IA32_RTIT_ADDRn_[AB], see. Section 35.2.4.3. The IP payload will be the target of the branch.</li> <li>▪ ContextEn: This is set on a CPL change, a CR3 write or any other means of changing ContextEn. The IP payload will be the Next IP of the instruction that changes context if it is not a branch, otherwise it will be the target of the branch.</li> </ul> |                 |   |                     |  |   |   |   |   |
| Application   | TIP.PGE packets bind to the instruction at the IP given in the payload.  |                 |   |                     |  |   |   |   |   |

### 35.4.2.5 Packet Generation Disable (TIP.PGD) Packet

**Table 35-21. TIP.PGD Packet Definition**

|               |   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
|---------------|---|---------------------|---|---|---|---|---|---|---|---|---|---|---|---------|--|--|---|---|---|---|---|---|---------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|---|-----------------|--|--|--|--|--|--|--|
| Name          | Target IP - Packet Generation Disable (TIP.PGD) Packet  |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Packet Format | <table border="1" data-bbox="318 380 1305 751"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td colspan="3">IPBytes</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="8">TargetIP[7:0]</td> </tr> <tr> <td>2</td> <td colspan="8">TargetIP[15:8]</td> </tr> <tr> <td>3</td> <td colspan="8">TargetIP[23:16]</td> </tr> <tr> <td>4</td> <td colspan="8">TargetIP[31:24]</td> </tr> <tr> <td>5</td> <td colspan="8">TargetIP[39:32]</td> </tr> <tr> <td>6</td> <td colspan="8">TargetIP[47:40]</td> </tr> <tr> <td>7</td> <td colspan="8">TargetIP[55:48]</td> </tr> <tr> <td>8</td> <td colspan="8">TargetIP[63:56]</td> </tr> </table>  |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | IPBytes |  |  | 0 | 0 | 0 | 0 | 1 | 1 | TargetIP[7:0] |  |  |  |  |  |  |  | 2 | TargetIP[15:8] |  |  |  |  |  |  |  | 3 | TargetIP[23:16] |  |  |  |  |  |  |  | 4 | TargetIP[31:24] |  |  |  |  |  |  |  | 5 | TargetIP[39:32] |  |  |  |  |  |  |  | 6 | TargetIP[47:40] |  |  |  |  |  |  |  | 7 | TargetIP[55:48] |  |  |  |  |  |  |  | 8 | TargetIP[63:56] |  |  |  |  |  |  |  |
|               | 7   | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 0             | IPBytes   |                     |   | 0 | 0 | 0 | 0 | 1 |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 1             | TargetIP[7:0]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 2             | TargetIP[15:8]  |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 3             | TargetIP[23:16]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 4             | TargetIP[31:24]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 5             | TargetIP[39:32]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 6             | TargetIP[47:40]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 7             | TargetIP[55:48]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| 8             | TargetIP[63:56]   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Dependencies  | PacketEn transitions to 0   | Generation Scenario | Any branch instruction, control flow transfer, or MOV CR3 that clears PacketEn, a WRMSR that disables packet generation and clears PacketEn |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Description   | <p>Indicates that PacketEn has transitioned to 0. It will include the IP at which the tracing ends, unless ContextEn = 0 or TraceEn=0 at the conclusion of the instruction or event that cleared PacketEn.</p> <p>PacketEn can be cleared due to any of the enables that comprise PacketEn transitioning from 1 to 0. Examples:</p> <ul style="list-style-type: none"> <li>▪ TriggerEn: This is cleared on software write to clear IA32_RTIT_CTL.TraceEn, or when IA32_RTIT_STATUS.Stopped is set, or on operational error. The IP payload will be suppressed in this case, and the “IPBytes” field will have the value 0.</li> <li>▪ FilterEn: This is cleared when software jumps out of the tracing region. This region is defined by enabling IP filtering in IA32_RTIT_CTL.ADDRn_CFG, and defining the range in IA32_RTIT_ADDRn_[AB], see. Section 35.2.4.3. The IP payload will depend on the type of the branch. For conditional branches, the payload is suppressed (IPBytes = 0), and in this case the destination can be inferred from the disassembly. For any other type of branch, the IP payload will be the target of the branch.</li> <li>▪ ContextEn: This can happen on a CPL change, a CR3 write or any other means of changing ContextEn. See Section 35.2.4.3 for details. In this case, when ContextEn is cleared, there will be no IP payload. The “IPBytes” field will have value 0.</li> </ul> <p>Note that, in cases where a branch that would normally produce a TIP packet (i.e., far transfer, indirect branch, interrupt, etc) or TNT update (conditional branch or compressed RT) causes PacketEn to transition from 1 to 0, the TIP or TNT bit will be replaced with TIP.PGD. The payload of the TIP.PGD will be the target of the branch, unless the result of the instruction causes TraceEn or ContextEn to be cleared (ie, SYSCALL when IA32_RTIT_CTL.OS=0, In the case where a conditional branch clears FilterEn and hence PacketEn, there will be no TNT bit for this branch, replaced instead by the TIP.PGD.</p> |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |
| Application   | <p>TIP.PGD can be produced by any branch instructions, as well as some non-branch instructions, that clear PacketEn. When produced by a branch, it replaces any TIP or TNT update that the branch would normally produce.</p> <p>In cases where there is an unbound FUP preceding the TIP.PGD, then the TIP.PGD is part of compound operation (i.e., asynchronous event or TSX abort) which cleared PacketEn. For most such cases, the TIP.PGD is simply replacing a TIP, and should be treated the same way. The TIP.PGD may or may not have an IP payload, depending on whether the operation cleared ContextEn.</p> <p>If there is not an associated FUP, the binding will depend on whether there is an IP payload. If there is an IP payload, then the TIP.PGD should be applied to either the next direct branch whose target matches the TIP.PGD payload, or the next branch that would normally generate a TIP or TNT packet. If there is no IP payload, then the TIP.PGD should apply to the next branch or MOV CR3 instruction.</p>   |                     |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |   |                 |  |  |  |  |  |  |  |

### 35.4.2.6 Flow Update (FUP) Packet

Table 35-22. FUP Packet Definition

| Name          | Flow Update (FUP) Packet  |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
|---------------|---|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---------|--|--|---|---|---|---|---|---|---------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td colspan="3">IPBytes</td> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="8">IP[7:0]</td> </tr> <tr> <td>2</td> <td colspan="8">IP[15:8]</td> </tr> <tr> <td>3</td> <td colspan="8">IP[23:16]</td> </tr> <tr> <td>4</td> <td colspan="8">IP[31:24]</td> </tr> <tr> <td>5</td> <td colspan="8">IP[39:32]</td> </tr> <tr> <td>6</td> <td colspan="8">IP[47:40]</td> </tr> <tr> <td>7</td> <td colspan="8">IP[55:48]</td> </tr> <tr> <td>8</td> <td colspan="8">IP[63:56]</td> </tr> </tbody> </table> |                     |   |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | IPBytes |  |  | 1 | 1 | 1 | 0 | 1 | 1 | IP[7:0] |  |  |  |  |  |  |  | 2 | IP[15:8] |  |  |  |  |  |  |  | 3 | IP[23:16] |  |  |  |  |  |  |  | 4 | IP[31:24] |  |  |  |  |  |  |  | 5 | IP[39:32] |  |  |  |  |  |  |  | 6 | IP[47:40] |  |  |  |  |  |  |  | 7 | IP[55:48] |  |  |  |  |  |  |  | 8 | IP[63:56] |  |  |  |  |  |  |  |
|               | 7   | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 0             | IPBytes   |                     |   | 1 | 1 | 1 | 0 | 1 |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 1             | IP[7:0]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 2             | IP[15:8]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 3             | IP[23:16]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 4             | IP[31:24]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 5             | IP[39:32]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 6             | IP[47:40]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 7             | IP[55:48]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| 8             | IP[63:56]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn & ContextEn.<br>(Typically depends on BranchEn and FilterEn as well, see Section 35.2.4 for details.)   | Generation Scenario | Asynchronous Events (interrupts, exceptions, INIT, SIPI, SMI, VM exit, #MC), XBEGIN, XEND, XABORT, XACQUIRE, XRELEASE, EENTER, EEXIT, ERESUME, EEE, AEX, <sup>1</sup> INT 0, INT 3, INT n, a WRMSR that disables packet generation. |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| Description   | Provides the source address for asynchronous events, and some other instructions. Is never sent alone, always sent with an associated TIP or MODE packet, and potentially others.   |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |
| Application   | <p>FUP packets provide the IP to which they bind. However, they are never standalone, but are coupled with other packets.</p> <p>In TSX cases, the FUP is immediately preceded by a MODE.TSX, which binds to the same IP. A TIP will follow only in the case of TSX aborts, see Section 35.4.2.8 for details.</p> <p>Otherwise, FUPs are part of compound packet events (see Section 35.4.1). In these compound cases, the FUP provides the source IP for an instruction or event, while a following TIP (or TIP.PGD) packet will provide the destination IP. Other packets may be included in the compound event between the FUP and TIP.</p>  |                     |   |   |   |   |   |   |   |   |   |   |   |   |         |  |  |   |   |   |   |   |   |         |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |

**NOTES:**

1. EENTER, EEXIT, ERESUME, EEE, AEX apply only if Intel Software Guard Extensions is supported.

## FUP IP Payload

Flow Update Packet gives the source address of an instruction when it is needed. In general, branch instructions do not need a FUP, because the source address is clear from the disassembly. For asynchronous events, however, the source address cannot be inferred from the source, and hence a FUP will be sent. Table 35-23 illustrates cases where FUPs are sent, and which IP can be expected in those cases.

**Table 35-23. FUP Cases and IP Payload**

| Event  | Flow Update IP   | Comment  |
|--|--|--|
| External Interrupt, NMI/SMI, Traps, Machine Check (trap-like), INIT/SIPI | Address of next instruction (Next IP) that would have been executed    | Functionally, this matches the LBR FROM field value and also the EIP value which is saved onto the stack.  |
| Exceptions/Faults, Machine check (fault-like)                            | Address of the instruction which took the exception/fault (Current IP) | This matches the similar functionality of LBR FROM field value and also the EIP value which is saved onto the stack.   |
| Software Interrupt   | Address of the software interrupt instruction (Current IP)             | This matches the similar functionality of LBR FROM field value, but does not match the EIP value which is saved onto the stack (Next Linear Instruction Pointer - NLIP). |
| EENTER, EEXIT, ERESUME, Enclave Exiting Event (EEE), AEX <sup>1</sup>    | Current IP of the instruction  | This matches the LBR FROM field value and also the EIP value which is saved onto the stack.  |
| XACQUIRE   | Address of the X* instruction  |  |
| XRELEASE, XBEGIN, XEND, XABORT, other transactional abort                | Current IP   |  |
| #SMI   | IP that is saved into SMRAM  |  |
| WRMSR that clears TraceEn  | Current IP   |  |

### NOTES:

1. Information on EENTER, EEXIT, ERESUME, EEE, Asynchronous Enclave eXit (AEX) can be found in *Intel® Software Guard Extensions Programming Reference*.

On a canonical fault due to sequentially fetching an instruction in non-canonical space (as opposed to jumping to non-canonical space), the IP of the fault (and thus the payload of the FUP) will be a non-canonical address. This is consistent with what is pushed on the stack for such faulting cases.

If there are post-commit task switch faults, the IP value of the FUP will be the original IP when the task switch started. This is the same value as would be seen in the LBR\_FROM field. But it is a different value as is saved on the stack or VMCS.

### 35.4.2.7 Paging Information (PIP) Packet

Table 35-24. PIP Packet Definition

| Name          | Paging Information (PIP) Packet   |                |   |                     |   |   |   |   |         |  |
|---------------|---|----------------|---|---------------------|---|---|---|---|---------|--|
| Packet Format |   | 7              | 6 | 5                   | 4 | 3   | 2 | 1 | 0       |  |
|               | 0   | 0              | 0 | 0                   | 0 | 0   | 0 | 1 | 0       |  |
|               | 1   | 0              | 1 | 0                   | 0 | 0   | 0 | 1 | 1       |  |
|               | 2   | CR3[11:5] or 0 |   |                     |   |   |   |   | RSVD/NR |  |
|               | 3   | CR3[19:12]     |   |                     |   |   |   |   |         |  |
|               | 4   | CR3[27:20]     |   |                     |   |   |   |   |         |  |
|               | 5   | CR3[35:28]     |   |                     |   |   |   |   |         |  |
|               | 6   | CR3[43:36]     |   |                     |   |   |   |   |         |  |
|               | 7   | CR3[51:44]     |   |                     |   |   |   |   |         |  |
| Dependencies  | TriggerEn && ContextEn && IA32_RTIT_CTL.OS  |                |   | Generation Scenario |   | MOV CR3, Task switch, INIT, SIPI, PSB+, VM exit, VM entry |   |   |         |  |
| Description   | <p>The CR3 payload shown includes only the address portion of the CR3 value. For PAE paging, CR3[11:5] are thus included. For other paging modes (32-bit and 4-level paging<sup>1</sup>), these bits are 0.</p> <p>This packet holds the CR3 address value. It will be generated on operations that modify CR3:</p> <ul style="list-style-type: none"> <li>▪ MOV CR3 operation</li> <li>▪ Task Switch</li> <li>▪ INIT and SIPI</li> <li>▪ VM exit, if “conceal VMX from PT” VM-exit control is 0 (see Section 35.5.1)</li> <li>▪ VM entry, if “conceal VMX from PT” VM-entry control is 0</li> </ul> <p>PIPs are not generated, despite changes to CR3, on SMI and RSM. This is due to the special behavior on these operations, see Section 35.2.8.3 for details. Note that, for some cases of task switch where CR3 is not modified, no PIP will be produced.</p> <p>The purpose of the PIP is to indicate to the decoder which application is running, so that it can apply the proper binaries to the linear addresses that are being traced.</p> <p>The PIP packet contains the new CR3 value when CR3 is written.</p> <p>PIPs generated by VM entries set the NR bit. PIPs generated in VMX non-root operation set the NR bit if the “conceal VMX from PT” VM-execution control is 0 (see Section 35.5.1). All other PIPs clear the NR bit.</p> |                |   |                     |   |   |   |   |         |  |
| Application   | <p>The purpose of the PIP packet is to help the decoder uniquely identify what software is running at any given time. When a PIP is encountered, a decoder should do the following:</p> <ol style="list-style-type: none"> <li>1) If there was a prior unbound FUP (that is, a FUP not preceded by a packet such as MODE.TSX that consumes it, and it hence pairs with a TIP that has not yet been seen), then this PIP is part of a compound packet event (Section 35.4.1). Find the ending TIP and apply the new CR3/NR values to the TIP payload IP.</li> <li>2) Otherwise, look for the next MOV CR3, far branch, or VMRESUME/VMLAUNCH in the disassembly, and apply the new CR3 to the next (or target) IP.</li> </ol> <p>For examples of the packets generated by these flows, see Section 35.7.</p>  |                |   |                     |   |   |   |   |         |  |

#### NOTES:

1. Earlier versions of this manual used the term “IA-32e paging” to identify 4-level paging.

### 35.4.2.8 MODE Packets

MODE packets keep the decoder informed of various processor modes about which it needs to know in order to properly manage the packet output, or to properly disassemble the associated binaries. MODE packets include a header and a mode byte, as shown below.

**Table 35-25. General Form of MODE Packets**

|   |         |   |   |      |   |   |   |   |
|---|---------|---|---|------|---|---|---|---|
|   | 7       | 6 | 5 | 4    | 3 | 2 | 1 | 0 |
| 0 | 1       | 0 | 0 | 1    | 1 | 0 | 0 | 1 |
| 1 | Leaf ID |   |   | Mode |   |   |   |   |

The MODE Leaf ID indicates which set of mode bits are held in the lower bits.

#### MODE.Exec Packet

**Table 35-26. MODE.Exec Packet Definition**

| Name          | MODE.Exec Packet  |                     |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
|---------------|---|---------------------|---|------|------------------------|-----------------|------|--------------|-----|---|---|-------------|---|---|-------------|---|---|-------------|---|---|---|---|---|---|---|---|---|---|------|--------------|
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>CS.D</td> <td>(CS.L &amp; LMA)</td> </tr> </table>  |                     |   |      | 7                      | 6               | 5    | 4            | 3   | 2 | 1 | 0           | 0 | 1 | 0           | 0 | 1 | 1           | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | CS.D | (CS.L & LMA) |
|               | 7   | 6                   | 5   | 4    | 3                      | 2               | 1    | 0            |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 0             | 1   | 0                   | 0   | 1    | 1                      | 0               | 0    | 1            |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 1             | 0   | 0                   | 0   | 0    | 0                      | 0               | CS.D | (CS.L & LMA) |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| Dependencies  | PacketEn  | Generation Scenario | Far branch, interrupt, exception, VM exit, and VM entry, if the mode changes. PSB+, and any scenario that can generate a TIP.PGE, such that the mode may have changed since the last MODE.Exec. |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| Description   | <p>Indicates whether software is in 16, 32, or 64-bit mode, by providing the CS.D and (CS.L &amp; IA32_EFER.LMA) values. Essential for the decoder to properly disassemble the associated binary.</p> <table border="1"> <thead> <tr> <th>CS.D</th> <th>(CS.L &amp; IA32_EFER.LMA)</th> <th>Addressing Mode</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1</td> <td>N/A</td> </tr> <tr> <td>0</td> <td>1</td> <td>64-bit mode</td> </tr> <tr> <td>1</td> <td>0</td> <td>32-bit mode</td> </tr> <tr> <td>0</td> <td>0</td> <td>16-bit mode</td> </tr> </tbody> </table> <p>MODE.Exec is sent at the time of a mode change, if PacketEn=1 at the time, or when tracing resumes, if necessary. In the former case, the MODE.Exec packet is generated along with other packets that result from the far transfer operation that changes the mode. In cases where the mode changes while PacketEn=0, the processor will send out a MODE.Exec along with the TIP.PGE when tracing resumes. The processor may opt to suppress the MODE.Exec when tracing resumes if the mode matches that from the last MODE.Exec packet, if there was no PSB in between.</p> |                     |   | CS.D | (CS.L & IA32_EFER.LMA) | Addressing Mode | 1    | 1            | N/A | 0 | 1 | 64-bit mode | 1 | 0 | 32-bit mode | 0 | 0 | 16-bit mode |   |   |   |   |   |   |   |   |   |   |      |              |
| CS.D          | (CS.L & IA32_EFER.LMA)  | Addressing Mode     |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 1             | 1   | N/A                 |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 0             | 1   | 64-bit mode         |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 1             | 0   | 32-bit mode         |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| 0             | 0   | 16-bit mode         |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |
| Application   | MODE.Exec always immediately precedes a TIP or TIP.PGE. The mode change applies to the IP address in the payload of the next TIP or TIP.PGE.  |                     |   |      |                        |                 |      |              |     |   |   |             |   |   |             |   |   |             |   |   |   |   |   |   |   |   |   |   |      |              |

## MODE.TSX Packet

Table 35-27. MODE.TSX Packet Definition

| Name          | MODE.TSX Packet  |   |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
|---------------|--|---|---|---|---|---|---------|------|---------|------|-------------|---|---|-----|---|---|--|---|---|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---------|------|
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>TXAbort</td> <td>InTX</td> </tr> </table>  |   |   |   |   |   |         |      |         | 7    | 6           | 5 | 4 | 3   | 2 | 1 | 0  | 0 | 1 | 0                   | 0 | 1 | 1   | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 0 | 0 | TXAbort | InTX |
|               | 7  | 6   | 5   | 4 | 3 | 2 | 1       | 0    |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 0             | 1  | 0   | 0   | 1 | 1 | 0 | 0       | 1    |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 1             | 0  | 0   | 1   | 0 | 0 | 0 | TXAbort | InTX |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| Dependencies  | TriggerEn and ContextEn  | Generation Scenario                                     | XBEGIN, XEND, XABORT, XACQUIRE, XRELEASE, if InTX changes, Asynchronous TSX Abort, PSB+ |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| Description   | <p>Indicates when a TSX transaction (either HLE or RTM) begins, commits, or aborts. Instructions executed transactionally will be “rolled back” if the transaction is aborted.</p> <table border="1"> <thead> <tr> <th>TXAbort</th> <th>InTX</th> <th>Implication</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1</td> <td>N/A</td> </tr> <tr> <td>0</td> <td>1</td> <td>Transaction begins, or executing transactionally</td> </tr> <tr> <td>1</td> <td>0</td> <td>Transaction aborted</td> </tr> <tr> <td>0</td> <td>0</td> <td>Transaction committed, or not executing transactionally</td> </tr> </tbody> </table> |   |   |   |   |   |         |      | TXAbort | InTX | Implication | 1 | 1 | N/A | 0 | 1 | Transaction begins, or executing transactionally | 1 | 0 | Transaction aborted | 0 | 0 | Transaction committed, or not executing transactionally |   |   |   |   |   |   |   |   |   |   |         |      |
| TXAbort       | InTX   | Implication   |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 1             | 1  | N/A   |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 0             | 1  | Transaction begins, or executing transactionally        |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 1             | 0  | Transaction aborted                                     |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| 0             | 0  | Transaction committed, or not executing transactionally |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |
| Application   | <p>If PacketEn=1, MODE.TSX always immediately precedes a FUP. If the TXAbort bit is zero, then the mode change applies to the IP address in the payload of the FUP. If TXAbort=1, then the FUP will be followed by a TIP, and the mode change will apply to the IP address in the payload of the TIP.</p> <p>MODE.TSX packets may be generated when PacketEn=0, due to FilterEn=0. In this case, only the last MODE.TSX generated before TIP.PGE need be applied.</p>  |   |   |   |   |   |         |      |         |      |             |   |   |     |   |   |  |   |   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |         |      |



### 35.4.2.9 TraceStop Packet

**Table 35-28. TraceStop Packet Definition**

|               |  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|---------------|--|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Name          | TraceStop Packet   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> </table>   |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
|               | 7  | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| 0             | 0  | 0                   | 0   | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| 1             | 1  | 0                   | 0   | 0 | 0 | 0 | 1 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Dependencies  | TriggerEn && ContextEn   | Generation Scenario | Taken branch with target in TraceStop IP region, MOV CR3 in TraceStop IP region, or WRMSR that sets TraceEn in TraceStop IP region. |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Description   | <p>Indicates when software has entered a user-configured TraceStop region. When the IP matches a TraceStop range while ContextEn and TriggerEn are set, a TraceStop action occurs. This disables tracing by setting IA32_RTIT_STATUS.Stopped, thereby clearing TriggerEn, and causes a TraceStop packet to be generated.</p> <p>The TraceStop action also forces FilterEn to 0. Note that TraceStop may not force a flush of internally buffered packets, and thus trace packet generation should still be manually disabled by clearing IA32_RTIT_CTL.TraceEn before examining output. See Section 35.2.4.3 for more details.</p> |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Application   | <p>If TraceStop follows a TIP.PGD (before the next TIP.PGE), then it was triggered either by the instruction that cleared PacketEn, or it was triggered by some later instruction that executed while FilterEn=0. In either case, the TraceStop can be applied at the IP of the TIP.PGD (if any).</p> <p>If TraceStop follows a TIP.PGE (before the next TIP.PGD), it should be applied at the last known IP.</p>  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |

### 35.4.2.10 Core:Bus Ratio (CBR) Packet

**Table 35-29. CBR Packet Definition**

|               |  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
|---------------|--|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----------------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|
| Name          | Core:Bus Ratio (CBR) Packet  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>2</td> <td colspan="8">Core:Bus Ratio</td> </tr> <tr> <td>3</td> <td colspan="8">Reserved</td> </tr> </table> |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 2 | Core:Bus Ratio |  |  |  |  |  |  |  | 3 | Reserved |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 0             | 0  | 0                   | 0   | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 1             | 0  | 0                   | 0   | 0 | 0 | 0 | 1 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 2             | Core:Bus Ratio   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 3             | Reserved   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn  | Generation Scenario | After any frequency change, on C-state wake up, PSB+, and after enabling trace packet generation. |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Description   | Indicates the core:bus ratio of the processor core. Useful for correlating wall-clock time and cycle time.   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Application   | All packets following the CBR represent instructions that executed with the new core:bus ratio, while all preceding packets (aside from timing packets) represent instructions that executed with the prior ratio. There is not a precise IP provided, to which to bind the CBR packet.  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |

## 35.4.2.11 Timestamp Counter (TSC) Packet

Table 35-30. TSC Packet Definition

| Name          | Timestamp Counter (TSC) Packet   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
|---------------|--|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|-------------|--|--|--|--|--|--|--|---|--------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="8">SW TSC[7:0]</td> </tr> <tr> <td>2</td> <td colspan="8">SW TSC[15:8]</td> </tr> <tr> <td>3</td> <td colspan="8">SW TSC[23:16]</td> </tr> <tr> <td>4</td> <td colspan="8">SW TSC[31:24]</td> </tr> <tr> <td>5</td> <td colspan="8">SW TSC[39:32]</td> </tr> <tr> <td>6</td> <td colspan="8">SW TSC[47:40]</td> </tr> <tr> <td>7</td> <td colspan="8">SW TSC[55:48]</td> </tr> </tbody> </table> |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | SW TSC[7:0] |  |  |  |  |  |  |  | 2 | SW TSC[15:8] |  |  |  |  |  |  |  | 3 | SW TSC[23:16] |  |  |  |  |  |  |  | 4 | SW TSC[31:24] |  |  |  |  |  |  |  | 5 | SW TSC[39:32] |  |  |  |  |  |  |  | 6 | SW TSC[47:40] |  |  |  |  |  |  |  | 7 | SW TSC[55:48] |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 0             | 0  | 0                   | 0   | 1 | 1 | 0 | 0 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 1             | SW TSC[7:0]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 2             | SW TSC[15:8]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 3             | SW TSC[23:16]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 4             | SW TSC[31:24]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 5             | SW TSC[39:32]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 6             | SW TSC[47:40]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| 7             | SW TSC[55:48]  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| Dependencies  | IA32_RTIT_CTL.TSCEn && TriggerEn   | Generation Scenario | Sent after any event that causes the processor clocks or Intel PT timing packets (such as MTC or CYC) to stop, This may include P-state changes, wake from C-state, or clock modulation. Also on transition of TraceEn from 0 to 1. |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| Description   | When enabled by software, a TSC packet provides the lower 7 bytes of the current TSC value, as returned by the RDTSC instruction. This may be useful for tracking wall-clock time, and synchronizing the packets in the log with other timestamped logs.   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |
| Application   | TSC packet provides a wall-clock proxy of the event which generated it (packet generation enable, sleep state wake, etc). In all cases, TSC does not precisely indicate the time of any control flow packets; however, all preceding packets represent instructions that executed before the indicated TSC time, and all subsequent packets represent instructions that executed after it. There is not a precise IP to which to bind the TSC packet.  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |             |  |  |  |  |  |  |  |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |

### 35.4.2.12 Mini Time Counter (MTC) Packet

**Table 35-31. MTC Packet Definition**

|               |  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
|---------------|--|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|------------|--|--|--|--|--|--|--|
| Name          | Mini time Counter (MTC) Packet   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| Packet Format | <table border="1" style="width: 100%; text-align: center;"> <tr> <td style="width: 10%;"></td> <td style="width: 10%;">7</td> <td style="width: 10%;">6</td> <td style="width: 10%;">5</td> <td style="width: 10%;">4</td> <td style="width: 10%;">3</td> <td style="width: 10%;">2</td> <td style="width: 10%;">1</td> <td style="width: 10%;">0</td> </tr> <tr> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="8">CTC[N+7:N]</td> </tr> </table>  |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | CTC[N+7:N] |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| 0             | 0  | 1                   | 0   | 1 | 1 | 0 | 0 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| 1             | CTC[N+7:N]   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| Dependencies  | IA32_RTIT_CTL.MTCEn && TriggerEn   | Generation Scenario | Periodic, based on the core crystal clock, or Always Running Timer (ART). |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| Description   | <p>When enabled by software, an MTC packet provides a periodic indication of wall-clock time. The 8-bit CTC (Common Timestamp Copy) payload value is set to <math>(ART \gg N) \&amp; FFH</math>. The frequency of the ART is related to the Maximum Non-Turbo frequency, and the ratio can be determined from CPUID leaf 15H, as described in Section 35.8.3. Software can select the threshold N, which determines the MTC frequency by setting the IA32_RTIT_CTL.MTCFreq field (see Section 35.2.7.2) to a supported value using the lookup enumerated by CPUID (see Section 35.3.1). See Section 35.8.3 for details on how to use the MTC payload to track TSC time.</p> <p>MTC provides 8 bits from the ART, starting with the bit selected by MTCFreq to dictate the frequency of the packet. Whenever that 8-bit range being watched changes, an MTC packet will be sent out with the new value of that 8-bit range. This allows the decoder to keep track of how much wall-clock time has elapsed since the last TSC packet was sent, by keeping track of how many MTC packets were sent and what their value was. The decoder can infer the truncated bits, CTC[N-1:0], are 0 at the time of the MTC packet.</p> <p>There are cases in which MTC packet can be dropped, due to overflow or other micro-architectural conditions. The decoder should be able to recover from such cases by checking the 8-bit payload of the next MTC packet, to determine how many MTC packets were dropped. It is not expected that &gt;256 consecutive MTC packets should ever be dropped.</p> |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |
| Application   | MTC does not precisely indicate the time of any other packet, nor does it bind to any IP. However, all preceding packets represent instructions or events that executed before the indicated ART time, and all subsequent packets represent instructions that executed after, or at the same time as, the ART time.  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |            |  |  |  |  |  |  |  |

## 35.4.2.13 TSC/MTC Alignment (TMA) Packet

Table 35-32. TMA Packet Definition

| Name          | TSC/MTC Alignment (TMA) Packet  |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
|---------------|---|---------------------|---------------------------|---|---|---|---|-------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----------|--|--|--|--|--|--|--|---|-----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|---|---|------------------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|-------|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>0</td> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>2</td> <td colspan="8">CTC[7:0]</td> </tr> <tr> <td>3</td> <td colspan="8">CTC[15:8]</td> </tr> <tr> <td>4</td> <td colspan="7">Reserved</td> <td>0</td> </tr> <tr> <td>5</td> <td colspan="8">FastCounter[7:0]</td> </tr> <tr> <td>6</td> <td colspan="7">Reserved</td> <td>FC[8]</td> </tr> </tbody> </table> |                     |                           |   | 7 | 6 | 5 | 4     | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | 2 | CTC[7:0] |  |  |  |  |  |  |  | 3 | CTC[15:8] |  |  |  |  |  |  |  | 4 | Reserved |  |  |  |  |  |  | 0 | 5 | FastCounter[7:0] |  |  |  |  |  |  |  | 6 | Reserved |  |  |  |  |  |  | FC[8] |
|               | 7   | 6                   | 5                         | 4 | 3 | 2 | 1 | 0     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 0             | 0   | 0                   | 0                         | 0 | 0 | 0 | 1 | 0     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 1             | 0   | 1                   | 1                         | 1 | 0 | 0 | 1 | 1     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 2             | CTC[7:0]  |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 3             | CTC[15:8]   |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 4             | Reserved  |                     |                           |   |   |   |   | 0     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 5             | FastCounter[7:0]  |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| 6             | Reserved  |                     |                           |   |   |   |   | FC[8] |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| Dependencies  | IA32_RTIT_CTL.MTCEn &&<br>IA32_RTIT_CTL.TSCEn && TriggerEn  | Generation Scenario | Sent with any TSC packet. |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| Description   | The TMA packet serves to provide the information needed to allow the decoder to correlate MTC packets with TSC packets. With this packet, when a MTC packet is encountered, the decoder can determine how many timestamp counter ticks have passed since the last TSC or MTC packet. See Section 35.8.3.2 for details on how to make this calculation.  |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |
| Application   | TMA is always sent immediately following a TSC packet, and the payload values are consistent with the TSC payload value. Thus the application of TMA matches that of TSC.   |                     |                           |   |   |   |   |       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |          |  |  |  |  |  |  |  |   |           |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |       |

### 35.4.2.14 Cycle Count (CYC) Packet

**Table 35-33. Cycle Count Packet Definition**

| Name          | Cycle Count (CYC) Packet   |                     |  |   |   |     |   |     |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
|---------------|--|---------------------|--|---|---|-----|---|-----|---|---|---|---|---|--------------------|--|--|--|--|-----|---|---|---|---------------------|--|--|--|--|--|--|-----|---|----------------------|--|--|--|--|--|--|-----|-----|---------------------------------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%;"></th> <th style="width: 10%;">7</th> <th style="width: 10%;">6</th> <th style="width: 10%;">5</th> <th style="width: 10%;">4</th> <th style="width: 10%;">3</th> <th style="width: 10%;">2</th> <th style="width: 10%;">1</th> <th style="width: 10%;">0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td colspan="5">Cycle Counter[4:0]</td> <td>Exp</td> <td>1</td> <td>1</td> </tr> <tr> <td>1</td> <td colspan="7">Cycle Counter[11:5]</td> <td>Exp</td> </tr> <tr> <td>2</td> <td colspan="7">Cycle Counter[18:12]</td> <td>Exp</td> </tr> <tr> <td>...</td> <td colspan="8">... (if Exp = 1 in the previous byte)</td> </tr> </tbody> </table>  |                     |  |   | 7 | 6   | 5 | 4   | 3 | 2 | 1 | 0 | 0 | Cycle Counter[4:0] |  |  |  |  | Exp | 1 | 1 | 1 | Cycle Counter[11:5] |  |  |  |  |  |  | Exp | 2 | Cycle Counter[18:12] |  |  |  |  |  |  | Exp | ... | ... (if Exp = 1 in the previous byte) |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5  | 4 | 3 | 2   | 1 | 0   |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| 0             | Cycle Counter[4:0]   |                     |  |   |   | Exp | 1 | 1   |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| 1             | Cycle Counter[11:5]  |                     |  |   |   |     |   | Exp |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| 2             | Cycle Counter[18:12]   |                     |  |   |   |     |   | Exp |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| ...           | ... (if Exp = 1 in the previous byte)  |                     |  |   |   |     |   |     |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| Dependencies  | IA32_RTIT_CTL.CYCEn && TriggerEn   | Generation Scenario | Can be sent at any time, though a maximum of one CYC packet is sent per core clock cycle. See Section 35.3.6 for CYC-eligible packets. |   |   |     |   |     |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| Description   | <p>The Cycle Counter field increments at the same rate as the processor core clock ticks, but with a variable length format (using a trailing EXP bit field) and a range-capped byte length.</p> <p>If the CYC value is less than 32, a 1-byte CYC will be generated, with Exp=0. If the CYC value is between 32 and 4095 inclusive, a 2-byte CYC will be generated, with byte 0 Exp=1 and byte 1 Exp=0. And so on.</p> <p>CYC provides the number of core clocks that have passed since the last CYC packet. CYC can be configured to be sent in every cycle in which an eligible packet is generated, or software can opt to use a threshold to limit the number of CYC packets, at the expense of some precision. These settings are configured using the IA32_RTIT_CTL.CycThresh field (see Section 35.2.7.2). For details on Cycle-Accurate Mode, IPC calculation, etc, see Section 35.3.6.</p> <p>When CycThresh=0, and hence no threshold is in use, then a CYC packet will be generated in any cycle in which any CYC-eligible packet is generated. The CYC packet will precede the other packets generated in the cycle, and provides the precise cycle time of the packets that follow.</p> <p>In addition to these CYC packets generated with other packets, CYC packets can be sent stand-alone. These packets serve simply to update the decoder with the number of cycles passed, and are used to ensure that a wrap of the processor's internal cycle counter doesn't cause cycle information to be lost. These stand-alone CYC packets do not indicate the cycle time of any other packet or operation, and will be followed by another CYC packet before any other CYC-eligible packet is seen.</p> <p>When CycThresh&gt;0, CYC packets are generated only after a minimum number of cycles have passed since the last CYC packet. Once this threshold has passed, the behavior above resumes, where CYC will either be sent in the next cycle that produces other CYC-eligible packets, or could be sent stand-alone.</p> <p>When using CYC thresholds, only the cycle time of the operation (instruction or event) that generates the CYC packet is truly known. Other operations simply have their execution time bounded: they completed at or after the last CYC time, and before the next CYC time.</p> |                     |  |   |   |     |   |     |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |
| Application   | <p>CYC provides the offset cycle time (since the last CYC packet) for the CYC-eligible packet that follows. If another CYC is encountered before the next CYC-eligible packet, the cycle values should be accumulated and applied to the next CYC-eligible packet.</p> <p>If a CYC packet is generated by a TNT, note that the cycle time provided by the CYC packet applies to the first branch in the TNT packet.</p>  |                     |  |   |   |     |   |     |   |   |   |   |   |                    |  |  |  |  |     |   |   |   |                     |  |  |  |  |  |  |     |   |                      |  |  |  |  |  |  |     |     |                                       |  |  |  |  |  |  |  |

## 35.4.2.15 VMCS Packet

Table 35-34. VMCS Packet Definition

| Name          | VMCS Packet   |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
|---------------|---|---------------------|---|---|---|---|---|---|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----------------------|--|--|--|--|--|--|--|---|----------------------|--|--|--|--|--|--|--|---|----------------------|--|--|--|--|--|--|--|---|----------------------|--|--|--|--|--|--|--|---|----------------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>2</td> <td colspan="8">VMCS pointer [19:12]</td> </tr> <tr> <td>3</td> <td colspan="8">VMCS pointer [27:20]</td> </tr> <tr> <td>4</td> <td colspan="8">VMCS pointer [35:28]</td> </tr> <tr> <td>5</td> <td colspan="8">VMCS pointer [43:36]</td> </tr> <tr> <td>6</td> <td colspan="8">VMCS pointer [51:44]</td> </tr> </tbody> </table>  |                     |   |   |   |   |   |   |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 2 | VMCS pointer [19:12] |  |  |  |  |  |  |  | 3 | VMCS pointer [27:20] |  |  |  |  |  |  |  | 4 | VMCS pointer [35:28] |  |  |  |  |  |  |  | 5 | VMCS pointer [43:36] |  |  |  |  |  |  |  | 6 | VMCS pointer [51:44] |  |  |  |  |  |  |  |
|               | 7   | 6                   | 5   | 4 | 3 | 2 | 1 | 0 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 0             | 0   | 0                   | 0   | 0 | 0 | 0 | 1 | 0 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 1             | 1   | 1                   | 0   | 0 | 1 | 0 | 0 | 0 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 2             | VMCS pointer [19:12]  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 3             | VMCS pointer [27:20]  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 4             | VMCS pointer [35:28]  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 5             | VMCS pointer [43:36]  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| 6             | VMCS pointer [51:44]  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn && ContextEn;<br>Also in VMX operation.   | Generation Scenario | Generated on successful VMPTRLD, and optionally on SMM VM exits and VM entries that return from SMM (see Section 35.5). |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| Description   | <p>The VMCS packet provides a VMCS pointer for a decoder to determine the transition of code contexts:</p> <ul style="list-style-type: none"> <li>On a successful VMPTRLD (i.e., a VMPTRLD that doesn't fault, fail, or VM exit), the VMCS packet contains the logical processor's VMCS pointer established by VMPTRLD (for subsequent execution of a VM guest context).</li> <li>An SMM VM exit loads the logical processor's VMCS pointer with the SMM-transfer VMCS pointer. If the "conceal VMX from PT" VM-exit control is 0 (see Section 35.5.1), a VMCS packet provides this pointer. See Section 35.6 on tracing inside and outside STM.</li> <li>A VM entry that returns from SMM loads the logical processor's VMCS pointer from a field in the SMM-transfer VMCS. If the "conceal VMX from PT" VM-entry control is 0, a VMCS packet provides this pointer. Whether the VM entry is to VMX root operation or VMX non-root operation is indicated by the PIP.NR bit.</li> </ul> <p>A VMCS packet generated before a VMCS pointer has been loaded, or after the VMCS pointer has been cleared will set all 64 bits in the VMCS pointer field.</p> <p>VMCS packets will not be seen on processors with IA32_VMX_MISC[bit 14]=0, as these processors do not allow TraceEn to be set in VMX operation.</p> |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |
| Application   | <p>The purpose of the VMCS packet is to help the decoder uniquely identify changes in the executing software context in situations that CR3 may not be unique.</p> <p>When a VMCS packet is encountered, a decoder should do the following:</p> <ul style="list-style-type: none"> <li>If there was a prior unbound FUP (that is, a FUP not preceded by a packet such as MODE.TSX that consumes it, and it hence pairs with a TIP that has not yet been seen), then this VMCS is part of a compound packet event (Section 35.4.1). Find the ending TIP and apply the new VMCS base pointer value to the TIP payload IP.</li> <li>Otherwise, look for the next VMPTRLD, VMRESUME, or VMLAUNCH in the disassembly, and apply the new VMCS base pointer on the next VM entry.</li> </ul> <p>For examples of the packets generated by these flows, see Section 35.7.</p>  |                     |   |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |   |                      |  |  |  |  |  |  |  |

### 35.4.2.16 Overflow (OVF) Packet

**Table 35-35. OVF Packet Definition**

| Name          | Overflow (OVF) Packet  |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|---------------|--|---------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> </tbody> </table>   |                     |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 |
|               |  | 7                   | 6   | 5 | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|               | 0  | 0                   | 0   | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|               | 1  | 1                   | 1   | 1 | 1 | 0 | 0 | 1 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Dependencies  | TriggerEn  | Generation Scenario | On resolution of internal buffer overflow |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Description   | OVF simply indicates to the decoder that an internal buffer overflow occurred, and packets were likely lost. If BranchEN= 1, OVF is followed by a FUP or TIP.PGE which will provide the IP at which packet generation resumes. See Section 35.3.8.   |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Application   | When an OVF packet is encountered, the decoder should skip to the IP given in the subsequent FUP or TIP.PGE. The cycle counter for the CYC packet will be reset at the time the OVF packet is sent. Software should reset its call stack depth on overflow, since no RET compression is allowed across an overflow. Similarly, any IP compression that follows the OVF is guaranteed to use as a reference LastIP the IP payload of an IP packet that preceded the overflow. |                     |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |

### 35.4.2.17 Packet Stream Boundary (PSB) Packet

**Table 35-36. PSB Packet Definition**

| Name          | Packet Stream Boundary (PSB) Packet  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|---------------|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>2</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>3</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>4</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>5</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>6</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>7</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>8</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>9</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>10</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>11</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>12</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>13</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>14</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> <tr><td>15</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr> </tbody> </table> |   |   |   |   |   |   |   |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 3 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 5 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 6 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 7 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 8 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 9 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 11 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 12 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 13 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 14 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 15 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |
|               |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 0  | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 1  | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 2  | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 3  | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 4  | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 5  | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 6  | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 7  | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 8  | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 9  | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 10   | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 11   | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 12   | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
|               | 13   | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
| 14            | 0  | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |
| 15            | 1  | 0 | 0 | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |   |

**Table 35-36. PSB Packet Definition (Contd.)**

|              |   |                     |  |
|--------------|---|---------------------|--|
| Dependencies | TriggerEn   | Generation Scenario | Periodic, based on the number of output bytes generated while tracing. PSB is sent when IA32_RTIT_STATUS.PacketByteCnt=0, and each time it crosses the software selected threshold after that. May be sent for other micro-architectural conditions as well. |
| Description  | PSB is a unique pattern in the packet output log, and hence serves as a sync point for the decoder. It is a pattern that the decoder can search for in order to get aligned on packet boundaries. This packet is periodic, based on the number of output bytes, as indicated by IA32_RTIT_STATUS.PacketByteCnt. The period is chosen by software, via IA32_RTIT_CTL.PSBFreq (see Section 35.2.7.2). Note, however, that the PSB period is not precise, it simply reflects the average number of output bytes that should pass between PSBs. The processor will make a best effort to insert PSB as quickly after the selected threshold is reached as possible. The processor also may send extra PSB packets for some micro-architectural conditions.<br>PSB also serves as the leading packet for a set of “status-only” packets collectively known as PSB+ (Section 35.3.7). |                     |  |
| Application  | When a PSB is seen, the decoder should interpret all following packets as “status only”, until either a PSBEND or OVF packet is encountered. “Status only” implies that the binding and ordering rules to which these packets normally adhere are ignored, and the state they carry can instead be applied to the IP payload in the FUP packet that is included.  |                     |  |

### 35.4.2.18 PSBEND Packet

**Table 35-37. PSBEND Packet Definition**

| Name          | PSBEND Packet  |                     |  |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|---------------|--|---------------------|--|---|---|---|---|---|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <th>0</th> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <th>1</th> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> </tbody> </table> |                     |  |   |   |   |   |   |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 1 |
|               | 7  | 6                   | 5  | 4 | 3 | 2 | 1 | 0 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| 0             | 0  | 0                   | 0  | 0 | 0 | 0 | 1 | 0 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| 1             | 0  | 0                   | 1  | 0 | 0 | 0 | 1 | 1 |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Dependencies  | TriggerEn  | Generation Scenario | Always follows PSB packet, separated by PSB+ packets |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Description   | PSBEND is simply a terminator for the series of “status only” (PSB+) packets that follow PSB (Section 35.3.7).   |                     |  |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Application   | When a PSBEND packet is seen, the decoder should cease to treat packets as “status only”.  |                     |  |   |   |   |   |   |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |



### 35.4.2.19 Maintenance (MNT) Packet

**Table 35-38. MNT Packet Definition**

| Name          | Maintenance (MNT) Packet   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
|---------------|--|---------------------|--------------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|--------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|----|----------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>2</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>3</td> <td colspan="8">Payload[7:0]</td> </tr> <tr> <td>4</td> <td colspan="8">Payload[15:8]</td> </tr> <tr> <td>5</td> <td colspan="8">Payload[23:16]</td> </tr> <tr> <td>6</td> <td colspan="8">Payload[31:24]</td> </tr> <tr> <td>7</td> <td colspan="8">Payload[39:32]</td> </tr> <tr> <td>8</td> <td colspan="8">Payload[47:40]</td> </tr> <tr> <td>9</td> <td colspan="8">Payload[55:48]</td> </tr> <tr> <td>10</td> <td colspan="8">Payload[63:56]</td> </tr> </tbody> </table> |                     |                          |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 2 | 1 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 3 | Payload[7:0] |  |  |  |  |  |  |  | 4 | Payload[15:8] |  |  |  |  |  |  |  | 5 | Payload[23:16] |  |  |  |  |  |  |  | 6 | Payload[31:24] |  |  |  |  |  |  |  | 7 | Payload[39:32] |  |  |  |  |  |  |  | 8 | Payload[47:40] |  |  |  |  |  |  |  | 9 | Payload[55:48] |  |  |  |  |  |  |  | 10 | Payload[63:56] |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5                        | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 0             | 0  | 0                   | 0                        | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 1             | 1  | 1                   | 0                        | 0 | 0 | 0 | 1 | 1 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 2             | 1  | 0                   | 0                        | 0 | 1 | 0 | 0 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 3             | Payload[7:0]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 4             | Payload[15:8]  |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 5             | Payload[23:16]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 6             | Payload[31:24]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 7             | Payload[39:32]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 8             | Payload[47:40]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 9             | Payload[55:48]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| 10            | Payload[63:56]   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn  | Generation Scenario | Implementation specific. |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| Description   | This packet is generated by hardware, the payload meaning is model-specific.   |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |
| Application   | Unless a decoder has been extended for a particular family/model/stepping to interpret MNT packet payloads, this packet should simply be ignored. It does not bind to any IP.  |                     |                          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |    |                |  |  |  |  |  |  |  |

### 35.4.2.20 PAD Packet

**Table 35-39. PAD Packet Definition**

| Name          | PAD Packet   |                     |                         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|---------------|--|---------------------|-------------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> </tr> </tbody> </table> |                     |                         |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|               | 7  | 6                   | 5                       | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |
| 0             | 0  | 0                   | 0                       | 0 | 0 | 0 | 0 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Dependencies  | TriggerEn  | Generation Scenario | Implementation specific |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Description   | PAD is simply a NOP packet. Processor implementations may choose to add pad packets to improve packet alignment or for implementation-specific reasons.  |                     |                         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Application   | Ignore PAD packets.  |                     |                         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |

## 35.4.2.21 PTWRITE (PTW) Packet

Table 35-40. PTW Packet Definition

| Name          | PTW Packet   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
|---------------|--|---------------------|---------------------|---|---|---|---|---|--------------|------------------|-----|---|-----|---|-----|----------|-----|----------|---|---|---|---|---|---|---|---|---|----|--------------|--|---|---|---|---|---|---|--------------|--|--|--|--|--|--|--|---|---------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|---|----------------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>IP</td> <td colspan="2">PayloadBytes</td> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>2</td> <td colspan="8">Payload[7:0]</td> </tr> <tr> <td>3</td> <td colspan="8">Payload[15:8]</td> </tr> <tr> <td>4</td> <td colspan="8">Payload[23:16]</td> </tr> <tr> <td>5</td> <td colspan="8">Payload[31:24]</td> </tr> <tr> <td>6</td> <td colspan="8">Payload[39:32]</td> </tr> <tr> <td>7</td> <td colspan="8">Payload[47:40]</td> </tr> <tr> <td>8</td> <td colspan="8">Payload[55:48]</td> </tr> <tr> <td>9</td> <td colspan="8">Payload[63:56]</td> </tr> </table> |                     |                     |   |   |   |   |   |              | 7                | 6   | 5 | 4   | 3 | 2   | 1        | 0   | 0        | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | IP | PayloadBytes |  | 1 | 0 | 0 | 1 | 0 | 2 | Payload[7:0] |  |  |  |  |  |  |  | 3 | Payload[15:8] |  |  |  |  |  |  |  | 4 | Payload[23:16] |  |  |  |  |  |  |  | 5 | Payload[31:24] |  |  |  |  |  |  |  | 6 | Payload[39:32] |  |  |  |  |  |  |  | 7 | Payload[47:40] |  |  |  |  |  |  |  | 8 | Payload[55:48] |  |  |  |  |  |  |  | 9 | Payload[63:56] |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5                   | 4 | 3 | 2 | 1 | 0 |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 0             | 0  | 0                   | 0                   | 0 | 0 | 0 | 1 | 0 |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 1             | IP   | PayloadBytes        |                     | 1 | 0 | 0 | 1 | 0 |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 2             | Payload[7:0]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 3             | Payload[15:8]  |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 4             | Payload[23:16]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 5             | Payload[31:24]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 6             | Payload[39:32]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 7             | Payload[47:40]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 8             | Payload[55:48]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| 9             | Payload[63:56]   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
|               | <p>The PayloadBytes field indicates the number of bytes of payload that follow the header bytes. Encodings are as follows:</p> <table border="1"> <thead> <tr> <th>PayloadBytes</th> <th>Bytes of Payload</th> </tr> </thead> <tbody> <tr> <td>'00</td> <td>4</td> </tr> <tr> <td>'01</td> <td>8</td> </tr> <tr> <td>'10</td> <td>Reserved</td> </tr> <tr> <td>'11</td> <td>Reserved</td> </tr> </tbody> </table> <p>IP bit indicates if a FUP, whose payload will be the IP of the PTWRITE instruction, will follow.</p>  |                     |                     |   |   |   |   |   | PayloadBytes | Bytes of Payload | '00 | 4 | '01 | 8 | '10 | Reserved | '11 | Reserved |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| PayloadBytes  | Bytes of Payload   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| '00           | 4  |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| '01           | 8  |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| '10           | Reserved   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| '11           | Reserved   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn & ContextEn & FilterEn & PTWEn   | Generation Scenario | PTWRITE Instruction |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| Description   | <p>Contains the value held in the PTWRITE operand.<br/> This packet is CYC-eligible, and hence will generate a CYC packet if IA32_RTIT_CTL.CYCEn=1 and any CYC Threshold has been reached.</p>   |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |
| Application   | <p>Binds to the associated PTWRITE instruction. The IP of the PTWRITE will be provided by a following FUP, when PTW.IP=1.</p>  |                     |                     |   |   |   |   |   |              |                  |     |   |     |   |     |          |     |          |   |   |   |   |   |   |   |   |   |    |              |  |   |   |   |   |   |   |              |  |  |  |  |  |  |  |   |               |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |   |                |  |  |  |  |  |  |  |

### 35.4.2.22 Execution Stop (EXSTOP) Packet

**Table 35-41. EXSTOP Packet Definition**

|               |   |                     |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
|---------------|---|---------------------|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|
| Name          | EXSTOP Packet   |                     |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| Packet Format | <table border="1" style="margin-left: 20px;"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>IP</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> </table> <p>IP bit indicates if a FUP will follow.</p>  |                     |  |   | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | IP | 1 | 1 | 0 | 0 | 0 | 1 | 0 |
|               | 7   | 6                   | 5  | 4 | 3 | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| 0             | 0   | 0                   | 0  | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| 1             | IP  | 1                   | 1  | 0 | 0 | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| Dependencies  | TriggerEn & PwrEvtEn  | Generation Scenario | C-state entry, P-state change, or other processor clock power-down. Includes : <ul style="list-style-type: none"> <li>▪ Entry to C-state deeper than C0.0</li> <li>▪ TM1/2</li> <li>▪ STPCLK#</li> <li>▪ Frequency change due to IA32_CLOCK_MODULATION, Turbo</li> </ul> |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| Description   | This packet indicates that software execution has stopped due to processor clock powerdown. Later packets will indicate when execution resumes.<br>If EXSTOP is generated while ContextEn is set, the IP bit will be set, and EXSTOP will be followed by a FUP packet containing the IP at which execution stopped. More precisely, this will be the IP of the oldest instruction that has not yet completed.<br>This packet is CYC-eligible, and hence will generate a CYC packet if IA32_RTIT_CTL.CYCEn=1 and any CYC Threshold has been reached. |                     |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |
| Application   | If a FUP follows EXSTOP (hence IP bit set), the EXSTOP can be bound to the FUP IP. Otherwise the IP is not known. Time of powerdown can be inferred from the preceding CYC, if CYCEn=1. Combined with the TSC at the time of wake (if TSCEn=1), this can be used to determine the duration of the powerdown.  |                     |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |   |   |   |   |   |   |   |

## 35.4.2.23 MWAIT Packet

Table 35-42. MWAIT Packet Definition

| Name          | MWAIT Packet   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
|---------------|--|---------------------|--|---|---|---|----------|---|--|--|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|------------------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|----------|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <thead> <tr> <th></th> <th>7</th> <th>6</th> <th>5</th> <th>4</th> <th>3</th> <th>2</th> <th>1</th> <th>0</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>2</td> <td colspan="8">MWAIT Hints[7:0]</td> </tr> <tr> <td>3</td> <td colspan="8">Reserved</td> </tr> <tr> <td>4</td> <td colspan="8">Reserved</td> </tr> <tr> <td>5</td> <td colspan="8">Reserved</td> </tr> <tr> <td>6</td> <td colspan="6">Reserved</td> <td colspan="2">EXT[1:0]</td> </tr> <tr> <td>7</td> <td colspan="8">Reserved</td> </tr> <tr> <td>8</td> <td colspan="8">Reserved</td> </tr> <tr> <td>9</td> <td colspan="8">Reserved</td> </tr> </tbody> </table> |                     |  |   |   |   |          |   |  |  | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 0 | 2 | MWAIT Hints[7:0] |  |  |  |  |  |  |  | 3 | Reserved |  |  |  |  |  |  |  | 4 | Reserved |  |  |  |  |  |  |  | 5 | Reserved |  |  |  |  |  |  |  | 6 | Reserved |  |  |  |  |  | EXT[1:0] |  | 7 | Reserved |  |  |  |  |  |  |  | 8 | Reserved |  |  |  |  |  |  |  | 9 | Reserved |  |  |  |  |  |  |  |
|               | 7  | 6                   | 5  | 4 | 3 | 2 | 1        | 0 |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 0             | 0  | 0                   | 0  | 0 | 0 | 0 | 1        | 0 |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 1             | 1  | 1                   | 0  | 0 | 0 | 0 | 1        | 0 |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 2             | MWAIT Hints[7:0]   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 3             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 4             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 5             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 6             | Reserved   |                     |  |   |   |   | EXT[1:0] |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 7             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 8             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 9             | Reserved   |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn & PwrEvtEn & ContextEn   | Generation Scenario | MWAIT instruction, or I/O redirection to MWAIT, that complete without fault or VMexit. |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Description   | Indicates that an MWAIT operation to C-state deeper than C0.0 completed. The MWAIT hints and extensions passed in by software are exposed in the payload.<br>This packet is CYC-eligible, and hence will generate a CYC packet if IA32_RTIT_CTL.CYCEn=1 and any CYC Threshold has been reached.  |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Application   | The MWAIT packet should bind to the IP of the next FUP, which will be the IP of the instruction that caused the MWAIT. This FUP will be shared with EXSTOP.  |                     |  |   |   |   |          |   |  |  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                  |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |          |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |

### 35.4.2.24 Power Entry (PWRE) Packet

**Table 35-43. PWRE Packet Definition**

|               |   |                     |   |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
|---------------|---|---------------------|---|---|-----------------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|----|----------|--|--|--|--|--|--|---|-------------------------|--|--|--|-----------------------------|--|--|--|
| Name          | PWRE Packet   |                     |   |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>2</td> <td>HW</td> <td colspan="7">Reserved</td> </tr> <tr> <td>3</td> <td colspan="4">Resolved Thread C-State</td> <td colspan="4">Resolved Thread Sub C-State</td> </tr> </table>  |                     |   |   | 7                           | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | 2 | HW | Reserved |  |  |  |  |  |  | 3 | Resolved Thread C-State |  |  |  | Resolved Thread Sub C-State |  |  |  |
|               | 7   | 6                   | 5   | 4 | 3                           | 2 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| 0             | 0   | 0                   | 0   | 0 | 0                           | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| 1             | 0   | 0                   | 1   | 0 | 0                           | 0 | 1 | 0 |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| 2             | HW  | Reserved            |   |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| 3             | Resolved Thread C-State   |                     |   |   | Resolved Thread Sub C-State |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| Dependencies  | TriggerEn & PwrEvtEn  | Generation Scenario | Transition to a C-state deeper than C0.0. |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| Description   | <p>Indicates processor entry to the resolved thread C-state and sub C-state indicated. The processor will remain in this C-state until either another PWRE indicates the processor has moved to a C-state deeper than C0.0, or a PWRX packet indicates a return to C0.</p> <p>Note that some CPUs may allow MWAIT to request a deeper C-state than is supported by the core. These deeper C-states may have platform-level implications that differentiate them. However, the PWRE packet will provide only the resolved thread C-state, which will not exceed that supported by the core.</p> <p>If the C-state entry was initiated by hardware, rather than a direct software request (such as MWAIT, HLT, or shut-down), the HW bit will be set to indicate this. Hardware Duty Cycling (see Section 14.5, "Hardware Duty Cycling (HDC)" in the <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B</i>) is an example of such a case.</p> |                     |   |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |
| Application   | <p>When transitioning from C0.0 to a deeper C-state, the PWRE packet will be followed by an EXSTOP. If that EXSTOP packet has the IP bit set, then the following FUP will provide the IP at which the C-state entry occurred. Subsequent PWRE packets generated before the next PWRX should bind to the same IP.</p>  |                     |   |   |                             |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |    |          |  |  |  |  |  |  |   |                         |  |  |  |                             |  |  |  |

### 35.4.2.25 Power Exit (PWRX) Packet

Table 35-44. PWRX Packet Definition

| Name          | PWRX Packet  |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
|---------------|--|---|---|-----|----------------------|---------|---|-----------|--|---|----------|---|---|----------------------------|---|---|---------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|-------------------|--|--|--|----------------------|--|--|--|---|----------|--|--|--|-------------|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|---|----------|--|--|--|--|--|--|--|
| Packet Format | <table border="1"> <tr> <td></td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> </tr> <tr> <td>2</td> <td colspan="4">Last Core C-State</td> <td colspan="4">Deepest Core C-State</td> </tr> <tr> <td>3</td> <td colspan="4">Reserved</td> <td colspan="4">Wake Reason</td> </tr> <tr> <td>4</td> <td colspan="8">Reserved</td> </tr> <tr> <td>5</td> <td colspan="8">Reserved</td> </tr> <tr> <td>6</td> <td colspan="8">Reserved</td> </tr> </table>   |   |   |     | 7                    | 6       | 5 | 4         | 3  | 2 | 1        | 0 | 0 | 0                          | 0                                       | 0 | 0       | 0   | 0 | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | 2 | Last Core C-State |  |  |  | Deepest Core C-State |  |  |  | 3 | Reserved |  |  |  | Wake Reason |  |  |  | 4 | Reserved |  |  |  |  |  |  |  | 5 | Reserved |  |  |  |  |  |  |  | 6 | Reserved |  |  |  |  |  |  |  |
|               | 7  | 6   | 5   | 4   | 3                    | 2       | 1 | 0         |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 0             | 0  | 0   | 0   | 0   | 0                    | 0       | 1 | 0         |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 1             | 1  | 0   | 1   | 0   | 0                    | 0       | 1 | 0         |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 2             | Last Core C-State  |   |   |     | Deepest Core C-State |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 3             | Reserved   |   |   |     | Wake Reason          |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 4             | Reserved   |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 5             | Reserved   |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 6             | Reserved   |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Dependencies  | TriggerEn & PwrEvtEn   | Generation Scenario                                     | Transition from a C-state deeper than C0.0 to C0. |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Description   | <p>Indicates processor return to thread C0 from a C-state deeper than C0.0. The Last Core C-State field provides the MWAIT encoding for the core C-state at the time of the wake. The Deepest Core C-State provides the MWAIT encoding for the deepest core C-state achieved during the sleep session, or since leaving thread C0. MWAIT encodings for C-states can be found in Table 4-11 in the <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B</i>. Note that these values reflect only the core C-state, and hence will not exceed the maximum supported core C-state, even if deeper C-states can be requested. The Wake Reason field is one-hot, encoded as follows:</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Field</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Interrupt</td> <td>Wake due to external interrupt received.</td> </tr> <tr> <td>1</td> <td>Reserved</td> <td></td> </tr> <tr> <td>2</td> <td>Store to Monitored Address</td> <td>Wake due to store to monitored address.</td> </tr> <tr> <td>3</td> <td>HW Wake</td> <td>Wake due to hardware autonomous condition, such as HDC.</td> </tr> </tbody> </table> |   |   | Bit | Field                | Meaning | 0 | Interrupt | Wake due to external interrupt received. | 1 | Reserved |   | 2 | Store to Monitored Address | Wake due to store to monitored address. | 3 | HW Wake | Wake due to hardware autonomous condition, such as HDC. |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Bit           | Field  | Meaning   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 0             | Interrupt  | Wake due to external interrupt received.                |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 1             | Reserved   |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 2             | Store to Monitored Address   | Wake due to store to monitored address.                 |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| 3             | HW Wake  | Wake due to hardware autonomous condition, such as HDC. |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |
| Application   | PWRX will always apply to the same IP as the PWRE. The time of wake can be discerned from (optional) timing packets that precede PWRX.   |   |   |     |                      |         |   |           |  |   |          |   |   |                            |   |   |         |   |   |   |   |   |   |   |   |   |   |   |   |   |   |                   |  |  |  |                      |  |  |  |   |          |  |  |  |             |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |   |          |  |  |  |  |  |  |  |

## 35.5 TRACING IN VMX OPERATION

On processors that IA32\_VMX\_MISC[bit 14] reports 1, TraceEn can be set in VMX operation. A series of mechanisms exist to allow the VMM to configure tracing based on the desired trace domain, and on the consumer of the trace output. The VMM can configure specific VMX controls to control what virtualization-specific data are included within the trace packets (see Section 35.5.1 for details). The MSR-load areas used by VMX transitions can be employed by the VMM to restrict tracing to the desired context (see Section 35.5.2 for details). These configuration options are summarized in Table 35-45. Table 35-45 covers common Intel PT usages while SMIs are handled by the default SMM treatment. Tracing with SMM Transfer Monitor is described in Section 35.6.

**Table 35-45. Common Usages of Intel PT and VMX**

| Target Domain           | Output Consumer    | Virtualize Output                    | Configure VMX Controls           | TraceEN Configuration   | Save/Restore MSR states of Trace Configuration     |
|-------------------------|--------------------|--------------------------------------|----------------------------------|---|--|
| System-Wide (VMM + VMs) | Host               | N/A                                  | Default setting (no suppression) | WRMSR or XRSTORS by Host  | N/A  |
| VMM Only                | Intel PT Aware VMM | N/A                                  | Enable suppression               | Use VMX MSR-load areas to disable tracing in VM, enable tracing on VM exits | N/A  |
| VM Only                 | Intel PT Aware VMM | N/A                                  | Enable suppression               | Use VMX MSR-load areas to enable tracing in VM, disable tracing on VM exits | N/A  |
| Intel PT Aware Guest(s) | Per Guest          | VMM adds trace output virtualization | Enable suppression               | Use VMX MSR-load areas to enable tracing in VM, disable tracing on VM exits | VMM updates guest state on VM exits due to XRSTORS |

### 35.5.1 VMX-Specific Packets and VMCS Controls

In all of the usages of VMX and Intel PT, a decoder in the host or VMM context can identify the occurrences of VMX transitions with the aid of VMX-specific packets. There are two kinds of packets relevant to VMX:

- **VMCS packet.** The VMX transitions of individual VMs can be distinguished by a decoder using the VMCS-pointer field in a VMCS packet. A VMCS packet is sent on a successful execution of VMPTRLD, and its VMCS-pointer field stores the VMCS pointer loaded by that execution. See Section 35.4.2.15 for details.
- **The NR (non-root) bit in a PIP packet.** Normally, the NR bit is set in any PIP packet generated in VMX non-root operation. In addition, PIP packets are generated with each VM entry and VM exit. Thus a transition of the NR bit from 0 to 1 indicates the occurrence of a VM entry, and a transition of 1 to 0 indicates the occurrence of a VM exit.

There are VMX controls that a VMM can set to conceal some of this VMX-specific information (by suppressing its recording) and thereby prevent it from leaking across virtualization boundaries. There is one of these controls (each of which is called “conceal VMX from PT”) of each type of VMX control.

**Table 35-46. VMX Controls For Intel Processor Trace**

| Type of VMX Control                            | Bit Position <sup>1</sup> | Value | Behavior  |
|--|---------------------------|-------|---|
| Secondary processor-based VM-execution control | 19                        | 0     | Each PIP generated in VM non-root operation will set the NR bit. PSB+ in VMX non-root operation will include the VMCS packet, to ensure that the decoder knows which guest is currently in use. |
|  |                           | 1     | Each PIP generated in VMX non-root operation will clear the NR bit. PSB+ in VMX non-root operation will not include the VMCS packet.  |
| VM-exit control                                | 24                        | 0     | Each VM exit generates a PIP in which the NR bit is clear. In addition, SMM VM exits generate VMCS packets.   |
|  |                           | 1     | VM exits do not generate PIPs, and no VMCS packets are generated on SMM VM exits.   |
| VM-entry control                               | 17                        | 0     | Each VM entry generates a PIP in which the NR bit is set (except VM entries that return from SMM to VMX root operation). In addition, VM entries that return from SMM generate VMCS packets.    |
|  |                           | 1     | VM entries do not generate PIPs, and no VMCS packets are generated on VM entries that return from SMM.  |

**NOTES:**

1. These are the positions of the control bits in the relevant VMX control fields.

The 0-settings of these VMX controls enable all VMX-specific packet information. The scenarios that would use these default settings also do not require the VMM to use VMX MSR-load areas to enable and disable trace-packet generation across VMX transitions.

If IA32\_VMX\_MISC[bit 14] reports 0, the 1-settings of the VMX controls in Table 35-46 are not supported, and VM entry will fail on any attempt to set them.

## 35.5.2 Managing Trace Packet Generation Across VMX Transitions

In tracing scenarios that collect packets for both VMX root operation and VMX non-root operation, a host executive can manage the MSRs associated with trace packet generation directly. The states of these MSRs need not be modified using MSR load areas across VMX transitions.

For tracing scenarios that collect packets only within VMX root operation or only within VMX non-root operation, the VMM can use the MSR load areas to toggle IA32\_RTIT\_CTL.TraceEn.

### 35.5.2.1 System-Wide Tracing

When a host or VMM configures Intel PT to collect trace packets of the entire system, it can leave the relevant VMX controls clear to allow VMX-specific packets to provide information across VMX transitions. The VMX MSR-load areas need not be used to load Intel PT MSRs on VM exits or VM entries.

The decoder will desire to identify the occurrence of VMX transitions. The packets of interests to a decoder are shown in Table 35-47.

**Table 35-47. Packets on VMX Transitions (System-Wide Tracing)**

| Event    | Packets             | Description  |
|----------|---------------------|--|
| VM exit  | FUP(GuestIP)        | The FUP indicates at which point in the guest flow the VM exit occurred. This is important, since VM exit can be an asynchronous event. The IP will match that written into the VMCS.  |
|          | PIP(HostCR3, NR=0)  | The PIP packet provides the new host CR3 value, as well as indication that the logical processor is entering VMX root operation. This allows the decoder to identify the change of executing context from guest to host and load the appropriate set of binaries to continue decode.   |
|          | TIP(HostIP)         | The TIP indicates the destination IP, the IP of the first instruction to be executed in VMX root operation.<br>Note, this packet could be preceded by a MODE.Exec packet (Section 35.4.2.8). This is generated only in cases where CS.D or (CS.L & EFER.LMA) change during the transition.   |
| VM entry | PIP(GuestCR3, NR=1) | The PIP packet provides the new guest CR3 value, as well as indication that the logical processor is entering VMX non-root operation. This allows the decoder to identify the change of executing context from host to guest and load the appropriate set of binaries to continue decode.  |
|          | TIP(GuestIP)        | The TIP indicates the destination IP, the IP of the first instruction to be executed in VMX non-root operation. This should match the RIP loaded from the VMCS.<br>Note, this packet could be preceded by a MODE.Exec packet (Section 35.4.2.8). This is generated only in cases where CS.D or (CS.L & EFER.LMA) change during the transition. |

Since the VMX controls that suppress packet generation are cleared, a VMCS packet will be included in all PSB+ for this usage scenario. Additionally, VMPTRLD will generate such a packet. Thus the decoder can distinguish the execution context of different VMs.

When the host VMM configures a system to collect trace packets in this scenario, it should emulate CPUID to report CPUID.(EAX=07H, ECX=0):EBX[bit 26] as 0 to guests, indicating to guests that Intel PT is not available.

### VMX TSC Manipulation

The TSC packets generated while in VMX non-root operation will include any changes resulting from the use of a VMM's use of the TSC offsetting or TSC scaling VMX controls (see Chapter 25, "VMX Non-Root Operation"). In this system-wide usage model, the decoder may need to account for the effect of per-VM adjustments in the TSC



packets generated in VMX non-root operation and the absence of TSC adjustments in TSC packets generated in VMX root operation. The VMM can supply this information to the decoder.

### 35.5.2.2 Host-Only Tracing

When trace packets in VMX non-root operation are not desired, the VMM can use the VM-entry MSR-load area to load IA32\_RTIT\_CTL (clearing TraceEn) to disable trace-packet generation in guests, and use the VM-exit MSR-load area to load IA32\_RTIT\_CTL to set TraceEn.

When tracing only the host, the decoder does not need information about the guests, and the VMX controls for suppressing VMX-specific packets can be set to reduce the packets generated. VMCS packets will still be generated on execution of VMPTRLD and in PSB+ generated in the host, but these will be unused by the decoder.

The packets of interests to a decoder when trace packets are collected for host-only tracing are shown in Table 35-48.

**Table 35-48. Packets on VMX Transitions (Host-Only Tracing)**

| Event    | Packets         | Description  |
|----------|-----------------|--|
| VM exit  | TIP.PGE(HostIP) | The TIP.PGE indicates that trace packet generation is enabled and gives the IP of the first instruction to be executed in VMX root operation.<br>Note, this packet could be preceded by a MODE.Exec packet (Section 35.4.2.8). This is generated only in cases where CS.D or (CS.L & EFER.LMA) change during the transition. |
| VM entry | TIP.PGD()       | The TIP indicates that trace packet generation was disabled. This ensure that all buffered packets are flushed out.  |

### 35.5.2.3 Guest-Only Tracing

A VMM can configure trace-packet generation while in VMX non-root operation for guests executing normally. This is accomplished by utilizing the VMX MSR-load areas on VM exits and VM entries to limit trace-packet generation to the guest environment.

For this usage, the VM-entry MSR load area is programmed to enable trace packet generation; the VM-exit MSR load area is used to clear IA32\_RTIT\_CTL.TraceEn so as to disable trace-packet generation in the host. Further, if it is preferred that the guest packet stream contain no indication that execution was in VMX non-root operation, the VMM should set to 1 all the VMX controls enumerated in Table 35-46.

### 35.5.2.4 Virtualization of Guest Output Packet Streams

Each Intel PT aware guest OS can produce one or more output packet streams to destination addresses specified as guest physical address using by context-switching IA32\_RTIT\_OUTPUT\_BASE within the guest. The processor generates trace packets to the physical address specified in IA32\_RTIT\_OUTPUT\_BASE, and those specified in the ToPA tables. Thus, a VMM that supports Intel PT aware guest OS may wish to virtualize the output configurations of IA32\_RTIT\_OUTPUT\_BASE and ToPA for each trace configuration state of all the guests.

### 35.5.2.5 Emulation of Intel PT Traced State

If a VMM emulates an element of processor state by taking a VM exit on reads and/or writes to that piece of state, and the state element impacts Intel PT packet generation or values, it may be incumbent upon the VMM to insert or modify the output trace data.

If a VM exit is taken on a guest write to CR3 (including "MOV CR3" as well as task switches), the PIP packet normally generated on the CR3 write will be missing.

To avoid decoder confusion when the guest trace is decoded, the VMM should emulate the missing PIP by writing it into the guest output buffer. If the guest CR3 value is manipulated, the VMM may also need to manipulate the IA32\_RTIT\_CR3\_MATCH value, in order to ensure the trace behavior matches the guest's expectation.

Similarly, if a VMM emulates the TSC value by taking a VM exit on RDTSC, the TSC packets generated in the trace may mismatch the TSC values returned by the VMM on RDTSC. To ensure that the trace can be properly aligned

with software logs based on RDTSC, the VMM should either make corresponding modifications to the TSC packet values in the guest trace, or use mechanisms such as TSC offsetting or TSC scaling in place of exiting.

### 35.5.2.6 TSC Scaling

When TSC scaling is enabled for a guest using Intel PT, the VMM should ensure that the value of Maximum Non-Turbo Ratio[15:8] in MSR\_PLATFORM\_INFO (MSR 0CEH) and the TSC/"core crystal clock" ratio (EBX/EAX) in CPUID leaf 15H are set in a manner consistent with the resulting TSC rate that will be visible to the VM. This will allow the decoder to properly apply TSC packets, MTC packets (based on the core crystal clock or ART, whose frequency is indicated by CPUID leaf 15H), and CBR packets (which indicate the ratio of the processor frequency to the Max Non-Turbo frequency). Absent this, or separate indication of the scaling factor, the decoder will be unable to properly track time in the trace. See Section 35.8.3 for details on tracking time within an Intel PT trace.

### 35.5.2.7 Failed VM Entry

The packets generated by a failed VM entry depend both on the VMCS configuration, as well as on the type of failure. The results to expect are summarized in the table below. Note that packets in *italics* may or may not be generated, depending on implementation choice, and the point of failure.

**Table 35-49. Packets on a Failed VM Entry**

| Usage Model | Entry Configuration                          | Early Failure (fall through to next IP) | Late Failure (VM-exit like)  |
|-------------|--|---|--|
| System-Wide | No use of VM-entry MSR-load area             | TIP (NextIP)                            | PIP(Guest CR3, NR=1), TraceEn 0->1 Packets (See Section 35.2.7.3), PIP(HostCR3, NR=0), TIP(HostIP) |
| VMM Only    | VM-entry MSR-load area used to clear TraceEn | TIP (NextIP)                            | TraceEn 0->1 Packets (See Section 35.2.7.3), TIP(HostIP)   |
| VM Only     | VM-entry MSR-load area used to set TraceEn   | None                                    | None   |

### 35.5.2.8 VMX Abort

VMX abort conditions take the processor into a shutdown state. On a VM exit that leads to VMX abort, some packets (FUP, PIP) may be generated, but any expected TIP, TIP.PGE, or TIP.PGD may be dropped.

## 35.6 TRACING AND SMM TRANSFER MONITOR (STM)

The SMM-transfer monitor (STM) is a VMM that operates inside SMM while in VMX root operation. An STM operates in conjunction with an executive monitor. The latter operates outside SMM and in VMX root operation. Transitions from the executive monitor or its VMs to the STM are called SMM VM exits. The STM returns from SMM via a VM entry to the VM in VMX non-root operation or the executive monitor in VMX root operation.

Intel PT supports tracing in an STM similar to tracing support for VMX operation as described above in Section 35.5. As a result, on a SMM VM exit resulting from #SMI, TraceEn is not saved and then cleared. Software can save the state of the trace configuration MSRs and clear TraceEn using the MSR load/save lists.

## 35.7 PACKET GENERATION SCENARIOS

Table 35-50 and Table 35-52 illustrate the packets generated in various scenarios. In the heading row, PacketEn is abbreviated as PktEn, ContextEn as CntxEn. Note that this assumes that TraceEn=1 in IA32\_RTIT\_CTL, while TriggerEn=1 and Error=0 in IA32\_RTIT\_STATUS, unless otherwise specified. Entries that do not matter in packet generation are marked "D.C." Packets followed by a "?" imply that these packets depend on additional factors, which are listed in the "Other Dependencies" column.

In Table 35-50, PktEn is evaluated based on TiggerEn & ContextEn & FilterEn & BranchEn.

**Table 35-50. Packet Generation under Different Enable Conditions**

| Case | Operation  | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output  |
|------|--|--------------|-------------|--------------|--|---|
| 1a   | Normal non-jump operation  | 0            | 0           | D.C.         |  | None  |
| 1b   | Normal non-jump operation  | 1            | 1           | 1            |  | None  |
| 2a   | WRMSR/XRSTORS/RSM that changes TraceEn 0 -> 1, with PacketByteCnt >0 | 0            | 0           | D.C.         | *TSC if TSCEn=1;<br>*TMA if TSCEn=MTCEn=1  | TSC?, TMA?, CBR   |
| 2b   | WRMSR/XRSTORS/RSM that changes TraceEn 0 -> 1, with PacketByteCnt =0 | 0            | 0           | D.C.         | *TSC if TSCEn=1;<br>*TMA if TSCEn=MTCEn=1  | PSB, PSBEND (see Section 35.4.2.17)   |
| 2d   | WRMSR/XRSTORS/RSM that changes TraceEn 0 -> 1, with PacketByteCnt >0 | 0            | 1           | 1            | TSC if TSCEn=1;<br>TMA if TSCEn=MTCEn=1  | TSC?, TMA?, CBR, MODE.Exec, TIP.PGE(NLIP)   |
| 2e   | WRMSR/XRSTORS/RSM that changes TraceEn 0 -> 1, with PacketByteCnt =0 | 0            | 1           | 1            |  | MODE.Exec, TIP.PGE(NLIP), PSB, PSBEND (see Section 35.4.2.8, 35.4.2.7, 35.4.2.13, 35.4.2.15, 35.4.2.17) |
| 3a   | WRMSR that changes TraceEn 1 -> 0                                    | 0            | 0           | D.C.         |  | None  |
| 3b   | WRMSR that changes TraceEn 1 -> 0                                    | 1            | 0           | D.C.         |  | FUP(CLIP), TIP.PGD()  |
| 5a   | MOV to CR3   | 0            | 0           | 0            |  | None  |
| 5f   | MOV to CR3   | 0            | 0           | 1            | TraceStop if executed in a TraceStop region  | PIP(NewCR3,NR?), TraceStop?   |
| 5b   | MOV to CR3   | 0            | 1           | 1            | *PIP.NR=1 if not in root operation and the "conceal VMX from PT" VM-execution control is 0<br>*MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | PIP(NewCR3, NR?), MODE.Exec?, TIP.PGE(NLIP)   |
| 5c   | MOV to CR3   | 1            | 0           | 0            |  | TIP.PGD()   |
| 5e   | MOV to CR3   | 1            | 0           | 1            | *PIP.NR=1 if not in root operation and the "conceal VMX from PT" VM-execution control is 0<br>*TraceStop if executed in a TraceStop region   | PIP(NewCR3, NR?), TIP.PGD(NLIP), TraceStop?   |
| 5d   | MOV to CR3   | 1            | 1           | 1            | *PIP.NR=1 if not in root operation and the "conceal VMX from PT" VM-execution control is 0   | PIP(NewCR3, NR?)  |
| 6a   | Unconditional direct near jump                                       | 0            | 0           | D.C.         |  | None  |
| 6b   | Unconditional direct near jump                                       | 1            | 0           | 1            | TraceStop if BLIP is in a TraceStop region   | TIP.PGD(BLIP), TraceStop?   |
| 6c   | Unconditional direct near jump                                       | 0            | 1           | 1            | MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB  | MODE.Exec?, TIP.PGE(BLIP)   |

Table 35-50. Packet Generation under Different Enable Conditions (Contd.)

| Case | Operation  | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output                               |
|------|--|--------------|-------------|--------------|--|--|
| 6d   | Unconditional direct near jump   | 1            | 1           | 1            |  | None   |
| 7a   | Conditional taken jump or compressed RET that does not fill up the internal TNT buffer | 0            | 0           | D.C.         |  | None   |
| 7b   | Conditional taken jump or compressed RET   | 0            | 1           | 1            | MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB  | MODE.Exec?, TIP.PGE(BLIP)                    |
| 7e   | Conditional taken jump or compressed RET, with empty TNT buffer                        | 1            | 0           | 1            | TraceStop if BLIP is in a TraceStop region   | TIP.PGD(), TraceStop?                        |
| 7f   | Conditional taken jump or compressed RET, with non-empty TNT buffer                    | 1            | 0           | 1            | TraceStop if BLIP is in a TraceStop region   | TNT, TIP.PGD(), TraceStop?                   |
| 7d   | Conditional taken jump or compressed RET that fills up the internal TNT buffer         | 1            | 1           | 1            |  | TNT  |
| 8a   | Conditional non-taken jump   | 0            | 0           | D.C.         |  | None   |
| 8d   | Conditional not-taken jump that fills up the internal TNT buffer                       | 1            | 1           | 1            |  | TNT  |
| 9a   | Near indirect jump (JMP, CALL, or uncompressed RET)                                    | 0            | 0           | D.C.         |  | None   |
| 9b   | Near indirect jump (JMP, CALL, or uncompressed RET)                                    | 0            | 1           | 1            | MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB  | MODE.Exec?, TIP.PGE(BLIP)                    |
| 9c   | Near indirect jump (JMP, CALL, or uncompressed RET)                                    | 1            | 0           | 1            | TraceStop if BLIP is in a TraceStop region   | TIP.PGD(BLIP), TraceStop?                    |
| 9d   | Near indirect jump (JMP, CALL, or uncompressed RET)                                    | 1            | 1           | 1            |  | TIP(BLIP)                                    |
| 10a  | Far Branch (CALL/JMP/RET)  | 0            | 0           | 0            |  | None   |
| 10f  | Far Branch (CALL/JMP/RET)  | 0            | 0           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1;<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region  | PIP(new CR3, NR?), TraceStop?                |
| 10b  | Far Branch (CALL/JMP/RET)  | 0            | 1           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1;<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | PIP(new CR3, NR?), MODE.Exec?, TIP.PGE(BLIP) |
| 10c  | Far Branch (CALL/JMP/RET)  | 1            | 0           | 0            |  | TIP.PGD()                                    |

**Table 35-50. Packet Generation under Different Enable Conditions (Contd.)**

| Case | Operation                 | PktEn Before | PktEn After | CntxEn After | Other Dependencies  | Packets Output  |
|------|---------------------------|--------------|-------------|--------------|---|---|
| 10d  | Far Branch (CALL/JMP/RET) | 1            | 0           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1;<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region   | PIP(new CR3, NR?), TIP.PGD(BLIP), TraceStop?            |
| 10e  | Far Branch (CALL/JMP/RET) | 1            | 1           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>* MODE.Exec if the operation changes CS.L/D or IA32_EFER.LMA                                     | PIP(NewCR3, NR?)?, MODE.Exec?, TIP(BLIP)                |
| 11a  | HW Interrupt              | 0            | 0           | 0            |   | None  |
| 11f  | HW Interrupt              | 0            | 0           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1;<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region   | PIP(new CR3, NR?), TraceStop?                           |
| 11b  | HW Interrupt              | 0            | 1           | 1            | *PIP if CR3 is updated (i.e., task switch), and OS=1;<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>* MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | PIP(new CR3, NR?), MODE.Exec?, TIP.PGE(BLIP)            |
| 11c  | HW Interrupt              | 1            | 0           | 0            |   | FUP(NLIP), TIP.PGD()                                    |
| 11d  | HW Interrupt              | 1            | 0           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region   | FUP(NLIP), PIP(NewCR3, NR?)?, TIP.PGD(BLIP), TraceStop? |

Table 35-50. Packet Generation under Different Enable Conditions (Contd.)

| Case | Operation       | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output  |
|------|-----------------|--------------|-------------|--------------|--|---|
| 11e  | HW Interrupt    | 1            | 1           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>* MODE.Exec if the operation changes CS.L/D or IA32_EFER.LMA                                   | FUP(NLIP), PIP(NewCR3, NR?)?, MODE.Exec?, TIP(BLIP)     |
| 12a  | SW Interrupt    | 0            | 0           | 0            |  | None  |
| 12f  | SW Interrupt    | 0            | 0           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region  | PIP(NewCR3, NR?)?, TraceStop?                           |
| 12b  | SW Interrupt    | 0            | 1           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | PIP(NewCR3, NR?)?, MODE.Exec?, TIP.PGE(BLIP)            |
| 12c  | SW Interrupt    | 1            | 0           | 0            |  | FUP(CLIP), TIP.PGD()                                    |
| 12d  | SW Interrupt    | 1            | 0           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region  | FUP(CLIP), PIP(NewCR3, NR?)?, TIP.PGD(BLIP), TraceStop? |
| 12e  | SW Interrupt    | 1            | 1           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>* MODE.Exec if the operation changes CS.L/D or IA32_EFER.LMA                                   | FUP(CLIP), PIP(NewCR3, NR?)?, MODE.Exec?, TIP(BLIP)     |
| 13a  | Exception/Fault | 0            | 0           | 0            |  | None  |

**Table 35-50. Packet Generation under Different Enable Conditions (Contd.)**

| Case | Operation                         | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output   |
|------|-----------------------------------|--------------|-------------|--------------|--|--|
| 13f  | Exception/Fault                   | 0            | 0           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region  | PIP(NewCR3, NR?)?, TraceStop?                                      |
| 13b  | Exception/Fault                   | 0            | 1           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | PIP(NewCR3, NR?)?, MODE.Exec?, TIP.PGE(BLIP)                       |
| 13c  | Exception/Fault                   | 1            | 0           | 0            |  | FUP(CLIP), TIP.PGD()   |
| 13d  | Exception/Fault                   | 1            | 0           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>*TraceStop if BLIP is in a TraceStop region  | FUP(CLIP), PIP(NewCR3, NR?)?, TIP.PGD(BLIP), TraceStop?            |
| 13e  | Exception/Fault                   | 1            | 1           | 1            | * PIP if CR3 is updated (i.e., task switch), and OS=1<br>*PIP.NR=1 if destination is not root operation and the "conceal VMX from PT" VM-execution control is 0;<br>* MODE.Exec if the operation changes CS.L/D or IA32_EFER.LMA                                   | FUP(CLIP), PIP(NewCR3, NR?)?, MODE.Exec?, TIP(BLIP)                |
| 14a  | SMI (TraceEn cleared)             | 0            | 0           | D.C.         |  | None   |
| 14b  | SMI (TraceEn cleared)             | 1            | 0           | 0            |  | FUP(SMRAM.LIP), TIP.PGD()  |
| 14f  | SMI (TraceEn cleared)             | 1            | 0           | 1            |  | NA   |
| 14c  | SMI (TraceEn cleared)             | 1            | 1           | 1            |  | NA   |
| 15a  | RSM, TraceEn restored to 0        | 0            | 0           | 0            |  | None   |
| 15b  | RSM, TraceEn restored to 1        | 0            | 0           | D.C.         |  | See WRMSR cases for packets on enable                              |
| 15c  | RSM, TraceEn restored to 1        | 0            | 1           | 1            |  | See WRMSR cases for packets on enable. FUP/TIP.PGE IP is SMRAM.LIP |
| 15e  | RSM (TraceEn=1, goes to shutdown) | 1            | 0           | 0            |  | None   |

Table 35-50. Packet Generation under Different Enable Conditions (Contd.)

| Case | Operation                              | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output  |
|------|--|--------------|-------------|--------------|--|---|
| 15f  | RSM (TraceEn=1, goes to shutdown)      | 1            | 0           | 1            |  | None  |
| 15d  | RSM (TraceEn=1, goes to shutdown)      | 1            | 1           | 1            |  | None  |
| 16i  | VM exit                                | 0            | 0           | 0            |  | None  |
| 16a  | VM exit                                | 0            | 0           | 1            | *PIP if OF=1 and the “conceal VMX from PT” VM-exit control is 0;<br>*TraceStop if VMCSH.LIP is in a TraceStop region         | PIP(HostCR3, NR=0)?, TraceStop?                                     |
| 16b  | VM exit, MSR list sets TraceEn=1       | 0            | 0           | 0            |  | See WRMSR cases for packets on enable. FUP IP is VMCSH.LIP          |
| 16c  | VM exit, MSR list sets TraceEn=1       | 0            | 1           | 1            |  | See WRMSR cases for packets on enable. FUP/TIP.PGE IP is VMCSH.LIP  |
| 16e  | VM exit                                | 0            | 1           | 1            | *PIP if OF=1 and the “conceal VMX from PT” VM-exit control is 0;<br>*MODE.Exec if the value is different, since last TIP.PGD | PIP(HostCR3, NR=0)?, MODE.Exec?, TIP.PGE(VMCSH.LIP)                 |
| 16f  | VM exit, MSR list clears TraceEn=0     | 1            | 0           | 0            | *PIP if OF=1 and the “conceal VMX from PT” VM-exit control is 0;   | FUP(VMCSG.LIP), PIP(HostCR3, NR=0)?, TIP.PGD                        |
| 16j  | VM exit, ContextEN 1->0                | 1            | 0           | 0            |  | FUP(VMCSG.LIP), TIP.PGD   |
| 16g  | VM exit                                | 1            | 0           | 1            | *PIP if OF=1 and the “conceal VMX from PT” VM-exit control is 0;<br>*TraceStop if VMCSH.LIP is in a TraceStop region         | FUP(VMCSG.LIP), PIP(HostCR3, NR=0)?, TIP.PGD(VMCSH.LIP), TraceStop? |
| 16h  | VM exit                                | 1            | 1           | 1            | *PIP if OF=1 and the “conceal VMX from PT” VM-exit control is 0;<br>*MODE.Exec if the value is different, since last TIP.PGD | FUP(VMCSG.LIP), PIP(HostCR3, NR=0)?, MODE.Exec, TIP(VMCSH.LIP)      |
| 17a  | VM entry                               | 0            | 0           | 0            |  | None  |
| 17b  | VM entry                               | 0            | 0           | 1            | *PIP if OF=1 and the “conceal VMX from PT” VM-entry control is 0;<br>*TraceStop if VMCSG.LIP is in a TraceStop region        | PIP(GuestCR3, NR=1)?, TraceStop?                                    |
| 17c  | VM entry, MSR load list sets TraceEn=1 | 0            | 0           | 1            |  | See WRMSR cases for packets on enable. FUP IP is VMCSG.LIP          |
| 17d  | VM entry, MSR load list sets TraceEn=1 | 0            | 1           | 1            |  | See WRMSR cases for packets on enable. FUP/TIP.PGE IP is VMCSG.LIP  |



Table 35-50. Packet Generation under Different Enable Conditions (Contd.)

| Case | Operation                           | PktEn Before | PktEn After | CntxEn After | Other Dependencies  | Packets Output                                       |
|------|-------------------------------------|--------------|-------------|--------------|---|--|
| 17f  | VM entry, FilterEN 0->1             | 0            | 1           | 1            | *PIP if OF=1 and the "conceal VMX from PT" VM-entry control is 0;<br>*MODE.Exec if the value is different, since last TIP.PGD | PIP(GuestCR3, NR=1)?, MODE.Exec?, TIP.PGE(VMCSg.LIP) |
| 17j  | VM entry, ContextEN 0->1            | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec, TIP.PGE(VMCSg.LIP)                        |
| 17g  | VM entry, MSR list clears TraceEn=0 | 1            | 0           | 0            | *PIP if OF=1 and the "conceal VMX from PT" VM-entry control is 0;   | PIP(GuestCR3, NR=1)?, TIP.PGD                        |
| 17h  | VM entry                            | 1            | 0           | 1            | *PIP if OF=1 and the "conceal VMX from PT" VM-entry control is 0;<br>*TraceStop if VMCSg.LIP is in a TraceStop region         | PIP(GuestCR3, NR=1)?, TIP.PGD(VMCSg.LIP), TraceStop? |
| 17i  | VM entry                            | 1            | 1           | 1            | *PIP if OF=1 and the "conceal VMX from PT" VM-entry control is 0;<br>*MODE.Exec if the value is different, since last TIP.PGD | PIP(GuestCR3, NR=1)?, MODE.Exec, TIP(VMCSg.LIP)      |
| 20a  | EENTER/ERESUME to non-debug enclave | 0            | 0           | 0            |   | None   |
| 20c  | EENTER/ERESUME to non-debug enclave | 1            | 0           | 0            |   | FUP(CLIP), TIP.PGD()                                 |
| 21a  | EEXIT from non-debug enclave        | 0            | 0           | D.C.         |   | None   |
| 21b  | EEXIT from non-debug enclave        | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec?, TIP.PGE(BLIP)                            |
| 22a  | AEX/EEE from non-debug enclave      | 0            | 0           | D.C.         |   | None   |
| 22b  | AEX/EEE from non-debug enclave      | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec?, TIP.PGE(AEP.LIP)                         |
| 23a  | EENTER/ERESUME to debug enclave     | 0            | 0           | D.C.         |   | None   |
| 23b  | EENTER/ERESUME to debug enclave     | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec?, TIP.PGE(BLIP)                            |
| 23c  | EENTER/ERESUME to debug enclave     | 1            | 0           | 0            |   | FUP(CLIP), TIP.PGD()                                 |
| 23d  | EENTER/ERESUME to debug enclave     | 0            | 0           | 1            | *TraceStop if BLIP is in a TraceStop region   | FUP(CLIP), TIP.PGD(BLIP), TraceStop?                 |
| 23e  | EENTER/ERESUME to debug enclave     | 1            | 1           | 1            |   | FUP(CLIP), TIP(BLIP)                                 |
| 24f  | EEXIT from debug enclave            | 0            | 0           | D.C.         |   | None   |
| 24b  | EEXIT from debug enclave            | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec?, TIP.PGE(BLIP)                            |
| 24d  | EEXIT from debug enclave            | 1            | 0           | 1            | *TraceStop if BLIP is in a TraceStop region   | FUP(CLIP), TIP.PGD(BLIP), TraceStop?                 |
| 24e  | EEXIT from debug enclave            | 1            | 1           | 1            |   | FUP(CLIP), TIP(BLIP)                                 |
| 25a  | AEX/EEE from debug enclave          | 0            | 0           | D.C.         |   | None   |

Table 35-50. Packet Generation under Different Enable Conditions (Contd.)

| Case | Operation                              | PktEn Before | PktEn After | CntxEn After | Other Dependencies   | Packets Output  |
|------|--|--------------|-------------|--------------|--|---|
| 25b  | AEX/EEE from debug enclave             | 0            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD   | MODE.Exec?, TIP.PGE(AEP.LIP)                            |
| 25d  | AEX/EEE from debug enclave             | 1            | 0           | 1            | *For AEX, FUP IP could be NLIP, for trap-like events   | FUP(CLIP), TIP.PGD(AEP.LIP)                             |
| 25e  | AEX/EEE from debug enclave             | 1            | 1           | 1            | *MODE.Exec if the value is different, since last TIP.PGD<br>*For AEX, FUP IP could be NLIP, for trap-like events | FUP(CLIP), MODE.Exec?, TIP(AEP.LIP)                     |
| 26a  | XBEGIN/XACQUIRE                        | 0            | 0           | D.C.         |  | None  |
| 26d  | XBEGIN/XACQUIRE that does not set InTX | 1            | 1           | 1            |  | None  |
| 26e  | XBEGIN/XACQUIRE that sets InTX         | 1            | 1           | 1            |  | MODE.TSX(InTX=1, TXAbort=0), FUP(CLIP)                  |
| 27a  | XEND/XRELEASE                          | 0            | 0           | D.C.         |  | None  |
| 27d  | XEND/XRELEASE that does not clear InTX | 1            | 1           | 1            |  | None  |
| 27e  | XEND/XRELEASE that clears InTX         | 1            | 1           | 1            |  | MODE.TSX(InTX=0, TXAbort=0), FUP(CLIP)                  |
| 28a  | XABORT(Async XAbort, or other)         | 0            | 0           | 0            |  | None  |
| 28e  | XABORT(Async XAbort, or other)         | 0            | 0           | 1            | *TraceStop if BLIP is in a TraceStop region  | MODE.TSX(InTX=0, TXAbort=1), TraceStop?                 |
| 28b  | XABORT(Async XAbort, or other)         | 0            | 1           | 1            |  | MODE.TSX(InTX=0, TXAbort=1), TIP.PGE(BLIP)              |
| 28c  | XABORT(Async XAbort, or other)         | 1            | 0           | 1            | *TraceStop if BLIP is in a TraceStop region  | MODE.TSX(InTX=0, TXAbort=1), TIP.PGD (BLIP), TraceStop? |
| 28d  | XABORT(Async XAbort, or other)         | 1            | 1           | 1            |  | MODE.TSX(InTX=0, TXAbort=1), FUP(CLIP), TIP(BLIP)       |
| 30a  | INIT (BSP)                             | 0            | 0           | 0            |  | None  |
| 30b  | INIT (BSP)                             | 0            | 0           | 1            | *TraceStop if RESET.LIP is in a TraceStop region   | PIP(0), TraceStop?                                      |
| 30c  | INIT (BSP)                             | 0            | 1           | 1            | * MODE.Exec if the value is different, since last TIP.PGD  | MODE.Exec?, PIP(0), TIP.PGE(ResetLIP)                   |
| 30d  | INIT (BSP)                             | 1            | 0           | 0            |  | FUP(NLIP), TIP.PGD()                                    |
| 30e  | INIT (BSP)                             | 1            | 0           | 1            | * PIP if OS=1<br>*TraceStop if RESET.LIP is in a TraceStop region  | FUP(NLIP), PIP(0), TIP.PGD, TraceStop?                  |
| 30f  | INIT (BSP)                             | 1            | 1           | 1            | * MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB<br>* PIP if OS=1 | FUP(NLIP), PIP(0)?, MODE.Exec?, TIP(ResetLIP)           |
| 31a  | INIT (AP, goes to wait-for-SIPI)       | 0            | D.C.        | D.C.         |  | None  |

**Table 35-50. Packet Generation under Different Enable Conditions (Contd.)**

| Case | Operation   | PktEn Before | PktEn After | CntxEn After | Other Dependencies  | Packets Output                |
|------|---|--------------|-------------|--------------|---|-------------------------------|
| 31b  | INIT (AP, goes to wait-for-SIPI)                        | 1            | D.C.        | D.C.         | * PIP if OS=1   | FUP(NLIP), PIP(0)             |
| 32a  | SIPI  | 0            | 0           | 0            |   | None                          |
| 32c  | SIPI  | 0            | 1           | 1            | * MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | MODE.Exec?, TIP.PGE(SIPI-LIP) |
| 32d  | SIPI  | 1            | 0           | 0            |   | TIP.PGD                       |
| 32e  | SIPI  | 1            | 0           | 1            | *TraceStop if SIPI LIP is in a TraceStop region   | TIP.PGD(SIPI LIP); TraceStop? |
| 32f  | SIPI  | 1            | 1           | 1            | * MODE.Exec if the mode has changed since the last MODE.Exec, or if no MODE.Exec since last PSB | MODE.Exec?, TIP(SIPI LIP)     |
| 33a  | MWAIT (to C0)   | D.C.         | D.C.        | D.C.         |   | None                          |
| 33b  | MWAIT (to higher-numbered C-State, packet sent on wake) | D.C.         | D.C.        | D.C.         | *TSC if TSCEn=1<br>*TMA if TSCEn=MTCEn=1  | TSC?, TMA?, CBR               |

In Table 35-52, PktEn is evaluated based on (TiggerEn & ContextEn & FilterEn & BranchEn & PwrEvtEn).

**Table 35-51. PwrEvtEn and PTWEn Packet Generation under Different Enable Conditions**

| Case  | Operation   | PktEn Before | PktEn After | CntxEn After | Other Dependencies                         | Packets Output  |
|-------|---|--------------|-------------|--------------|--|---|
| 16.1  | MWAIT or I/O redir to MWAIT, gets #UD or #GP fault  | D.C.         | D.C.        | D.C.         |  | None  |
| 16.2  | MWAIT or I/O redir to MWAIT, VM exits   | D.C.         | D.C.        | D.C.         |  | See VM exit examples (16[a-z] in Table 35-50) for BranchEn packets. |
| 16.3  | MWAIT or I/O redir to MWAIT, requests C0, or monitor not armed, or VMX virtual-interrupt delivery | D.C.         | D.C.        | D.C.         |  | None  |
| 16.4a | MWAIT(X) or I/O redir to MWAIT, goes to C-state Y (Y>0)   | D.C.         | 0           | 0            |  | PWRE(Cx), EXSTOP  |
| 16.4b | MWAIT(X) or I/O redir to MWAIT, goes to C-state Y (Y>0)   | D.C.         | D.C.        | 1            |  | MWAIT(Cy), PWRE(Cx), EXSTOP(IP), FUP(CLIP)                          |
| 16.5a | MWAIT(X) or I/O redir to MWAIT, Pending event after resolving to go to C-state Y (Y>0)            | D.C.         | 0           | 0            | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | PWRE(Cx), EXSTOP, TSC?, TMA?, CBR, PWRX(LCC, DCC, 0)                |
| 16.5b | MWAIT(X) or I/O redir to MWAIT, Pending event after resolving to go to C-state Y (Y>0)            | D.C.         | D.C.        | 1            | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | PWRE(Cx), EXSTOP(IP), FUP(CLIP), TSC?, TMA?, CBR, PWRX(LCC, DCC, 0) |
| 16.6a | MWAIT(5) or I/O redir to MWAIT, other thread(s) in core in C0/C1                                  | D.C.         | 0           | 0            |  | PWRE(C1), EXSTOP  |
| 16.6b | MWAIT(5) or I/O redir to MWAIT, other thread(s) in core in C0/C1                                  | D.C.         | D.C.        | 1            |  | MWAIT(5), PWRE(C1), EXSTOP(IP), FUP(CLIP)                           |

Table 35-51. PwrEvtEn and PTWEn Packet Generation under Different Enable Conditions (Contd.)

| Case   | Operation   | PktEn Before | PktEn After | CntxE After | Other Dependencies | Packets Output  |
|--------|---|--------------|-------------|-------------|--------------------|---|
| 16.9a  | HLT, Triple-fault shutdown, #MC with CR4.MCE=0, RSM to Cx (x>0) | D.C.         | 0           | 0           |                    | PWRE(C1), EXSTOP  |
| 16.9b  | HLT, Triple-fault shutdown, #MC with CR4.MCE=1, RSM to Cx (x>0) | D.C.         | D.C.        |             |                    | PWRE(C1), EXSTOP(IP), FUP(CLIP)   |
| 16.10a | VMX abort   | D.C.         | 0           | 0           |                    | See "VMX Abort" (cases 16* and 18* in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP                    |
| 16.10b | VMX abort   | D.C.         | D.C.        | 1           |                    | See "VMX Abort" (cases 16* and 18* in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP(IP), FUP(CLIP)     |
| 16.11a | RSM to Shutdown   | D.C.         | 0           | 0           |                    | See "RSM to Shutdown" (cases 15[def] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP                  |
| 16.11b | RSM to Shutdown   | D.C.         | D.C.        | 1           |                    | See "RSM to Shutdown" (cases 15[def] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP(IP), FUP(CLIP)   |
| 16.12a | INIT (BSP)  | D.C.         | 0           | 0           |                    | See "INIT (BSP)" (cases 30[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP                       |
| 16.12b | INIT (BSP)  | D.C.         | D.C.        | 1           |                    | See "INIT (BSP)" (cases 30[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP(IP), FUP(NLIP)        |
| 16.13a | INIT (AP, goes to Wait-for-SIPI)                                | D.C.         | 0           | 0           |                    | See "INIT (AP, goes to Wait-for-SIPI)" (cases 31[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP |

**Table 35-51. PwrEvtEn and PTWEn Packet Generation under Different Enable Conditions (Contd.)**

| Case   | Operation                           | PktEn Before | PktEn After | CntxE After | Other Dependencies                         | Packets Output   |
|--------|-------------------------------------|--------------|-------------|-------------|--|--|
| 16.13b | INIT (AP, goes to Wait-for-SIPI)    | D.C.         | D.C.        | 1           |  | See "INIT (AP, goes to Wait-for-SIPI)" (cases 31[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP(IP), FUP(NLIP) |
| 16.14a | Hardware Duty Cycling (HDC)         | D.C.         | 0           | 0           | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | PWRE(HW, C6), EXSTOP, TSC?, TMA?, CBR, PWRX(CC6, CC6, 0x8)   |
| 16.14b | Hardware Duty Cycling (HDC)         | D.C.         | D.C.        | 1           | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | PWRE(HW, C6), EXSTOP(IP), FUP(NLIP), TSC?, TMA?, CBR, PWRX(CC6, CC6, 0x8)  |
| 16.15a | VM entry to HLT or Shutdown         | D.C.         | 0           | 0           |  | See "VM entry" (cases 17[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP  |
| 16.15b | VM entry to HLT or Shutdown         | D.C.         | D.C.        | 1           |  | See "VM entry" (cases 17[a-z] in Table 35-50) for BranchEn packets that precede<br><br>PWRE(C1), EXSTOP(IP), FUP(CLIP)                         |
| 16.16a | EIST in C0, S1/TM1/TM2, or STP-CLK# | D.C.         | 0           | 0           | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | EXSTOP, TSC?, TMA?, CBR  |
| 16.16b | EIST in C0, S1/TM1/TM2, or STP-CLK# | D.C.         | D.C.        | 1           | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | EXSTOP(IP), FUP(NLIP), TSC?, TMA?, CBR   |
| 16.17  | EIST in Cx (x>0)                    | D.C.         | D.C.        | D.C.        |  | None   |
| 16.18  | INTR during Cx (x>0)                | D.C.         | D.C.        | D.C.        | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | TSC?, TMA?, CBR, PWRX(LCC, DCC, 0x1)<br><br>See "HW Interrupt" (cases 11[a-z] in Table 35-50) for BranchEn packets that follow.                |
| 16.18  | SMI during Cx (x>0)                 | D.C.         | D.C.        | D.C.        | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEn=1 | TSC?, TMA?, CBR, PWRX(LCC, DCC, 0)<br><br>See "HW Interrupt" (cases 14[a-z] in Table 35-50) for BranchEn packets that follow.                  |

**Table 35-51. PwrEvtEn and PTWEn Packet Generation under Different Enable Conditions (Contd.)**

| Case  | Operation  | PktEn Before | PktEn After | CntxEEn After | Other Dependencies                          | Packets Output   |
|-------|--|--------------|-------------|---------------|---|--|
| 16.19 | NMI during Cx (x>0)  | D.C.         | D.C.        | D.C.          | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEEn=1 | TSC?, TMA?, CBR,<br>PWRX(LCC, DCC, 0)<br><br>See “HW Interrupt” (cases 11[a-z] in Table 35-50) for BranchEn packets that follow. |
| 16.20 | Store to monitored address during Cx (x>0)   | D.C.         | D.C.        | D.C.          | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEEn=1 | TSC?, TMA?, CBR,<br>PWRX(LCC, DCC, 0x4)  |
| 16.22 | #MC, IERR, TSC deadline timer expiration, or APIC counter under-flow during Cx (x>0) | D.C.         | D.C.        | D.C.          | * TSC if TSCEn=1<br>* TMA if TSCEn=MTCEEn=1 | TSC?, TMA?, CBR,<br>PWRX(LCC, DCC, 0)  |

In Table 35-52, PktEn is evaluated based on (TiggerEn & ContextEn & FilterEn & BranchEn & PTWEn).

**Table 35-52. PwrEvtEn and PTWEn Packet Generation under Different Enable Conditions**

| Case   | Operation                 | PktEn Before | PktEn After | CntxEEn After | Other Dependencies        | Packets Output   |
|--------|---------------------------|--------------|-------------|---------------|---------------------------|--|
| 16.24a | PTWRITE rm32/64, no fault | D.C.         | D.C.        | D.C.          |                           | None   |
| 16.24b | PTWRITE rm32/64, no fault | D.C.         | 0           | 0             |                           | None   |
| 16.24d | PTWRITE rm32, no fault    | D.C.         | 1           | 1             | * FUP, IP=1 if FUPonPTW=1 | PTW(IP=1?, 4B,<br>rm32_value), FUP(CLIP)?                                  |
| 16.24e | PTWRITE rm64, no fault    | D.C.         | 1           | 1             | * FUP, IP=1 if FUPonPTW=1 | PTW(IP=1?, 8B,<br>rm64_value), FUP(CLIP)?                                  |
| 16.25a | PTWRITE mem32/64, fault   | D.C.         | D.C.        | D.C.          |                           | See “Exception/fault” (cases 13[a-z] in Table 35-50) for BranchEn packets. |

## 35.8 SOFTWARE CONSIDERATIONS

### 35.8.1 Tracing SMM Code

Nothing prevents an SMM handler from configuring and enabling packet generation for its own use. As described in Section 35.2.8.3, SMI will always clear TraceEn, so the SMM handler would have to set TraceEn in order to enable tracing. There are some unique aspects and guidelines involved with tracing SMM code, which follow:

1. SMM should save away the existing values of any configuration MSRs that SMM intends to modify for tracing. This will allow the non-SMM tracing context to be restored before RSM.
2. It is recommended that SMM wait until it sets CSbase to 0 before enabling packet generation, to avoid possible LIP vs RIP confusion.
3. Packet output cannot be directed to SMRR memory, even while tracing in SMM.
4. Before performing RSM, SMM should take care to restore modified configuration MSRs to the values they had immediately after #SMI. This involves first disabling packet generation by clearing TraceEn, then restoring any other configuration MSRs that were modified.
5. RSM

- Software must ensure that TraceEn=0 at the time of RSM. Tracing RSM is not a supported usage model, and the packets generated by RSM are undefined.
- For processors on which Intel PT and LBR use are mutually exclusive (see Section 35.3.1.2), any RSM during which TraceEn is restored to 1 will suspend any LBR or BTS logging.

## 35.8.2 Cooperative Transition of Multiple Trace Collection Agents

A third-party trace-collection tool should take into consideration the fact that it may be deployed on a processor that supports Intel PT but may run under any operating system.

In such a deployment scenario, Intel recommends that tool agents follow similar principles of cooperative transition of single-use hardware resources, similar to how performance monitoring tools handle performance monitoring hardware:

- Respect the “in-use” ownership of an agent who already configured the trace configuration MSRs, see architectural MSRs with the prefix “IA32\_RTIT\_” in Chapter 2, “Model-Specific Registers (MSRs)” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 4*, where “in-use” can be determined by reading the “enable bits” in the configuration MSRs.
- Relinquish ownership of the trace configuration MSRs by clearing the “enabled bits” of those configuration MSRs.

## 35.8.3 Tracking Time

This section describes the relationships of several clock counters whose update frequencies reside in different domains that feed into the timing packets. To track time, the decoder also needs to know the regularity or irregularity of the occurrences of various timing packets that store those clock counters.

Intel PT provides time information for three different but related domains:

- Processor timestamp counter

This counter increments at the max non-turbo or P1 frequency, and its value is returned on a RDTSC. Its frequency is fixed. The TSC packet holds the lower 7 bytes of the timestamp counter value. The TSC packet occurs occasionally and are much less frequent than the frequency of the time stamp counter. The timestamp counter will continue to increment when the processor is in deep C-States, with the exception of processors reporting CPUID.80000007H:EDX.InvariantTSC[bit 8] =0.

- Core crystal clock

The ratio of the core crystal clock to timestamp counter frequency is known as P, and can be calculated as CPUID.15H:EBX[31:0] / CPUID.15H:EAX[31:0]. The frequency of the core crystal clock is fixed and lower than that of the timestamp counter. The periodic MTC packet is generated based on software-selected multiples of the crystal clock frequency. The MTC packet is expected to occur more frequently than the TSC packet.

- Processor core clock

The processor core clock frequency can vary due to P-state and thermal conditions. The CYC packet provides elapsed time as measured in processor core clock cycles relative to the last CYC packet.

A decoder can use all or some combination of these packets to track time at different resolutions throughout the trace packets.

### 35.8.3.1 Time Domain Relationships

The three domains are related by the following formula:

$$\text{TimeStampValue} = (\text{CoreCrystalClockValue} * P) + \text{AdjustedProcessorCycles} + \text{Software\_Offset};$$

The CoreCrystalClockValue can provide the coarse-grained component of the TSC value. P, or the TSC/“core crystal clock” ratio, can be derived from CPUID leaf 15H, as described in Section 35.8.3.

The AdjustedProcessorCycles component provides the fine-grained distance from the rising edge of the last core crystal clock. Specifically, it is a cycle count in the same frequency as the timestamp counter from the last crystal clock rising edge. The value is adjusted based on the ratio of the processor core clock frequency to the Maximum Non-Turbo (or P1) frequency.

The Software\_Offsets component includes software offsets that are factored into the timestamp value, such as IA32\_TSC\_ADJUST.

### 35.8.3.2 Estimating TSC within Intel PT

For many usages, it may be useful to have an estimated timestamp value for all points in the trace. The formula provided in Section 35.8.3.1 above provides the framework for how such an estimate can be calculated from the various timing packets present in the trace.

The TSC packet provides the precise timestamp value at the time it is generated; however, TSC packets are infrequent, and estimates of the current timestamp value based purely on TSC packets are likely to be very inaccurate for this reason. In order to get more precise timing information between TSC packets, CYC packets and/or MTC packets should be enabled.

MTC packets provide incremental updates of the CoreCrystalClockValue. On processors that support CPUID leaf 15H, the frequency of the timestamp counter and the core crystal clock is fixed, thus MTC packets provide a means to update the running timestamp estimate. Between two MTC packets A and B, the number of crystal clock cycles passed is calculated from the 8-bit payloads of respective MTC packets:

$(CTC_B - CTC_A)$ , where  $CTC_i = MTC_i[15:8] \ll IA32\_RTIT\_CTL.MTCFreq$  and  $i = A, B$ .

The time from a TSC packet to the subsequent MTC packet can be calculated using the TMA packet that follows the TSC packet. The TMA packet provides both the crystal clock value (lower 16 bits, in the CTC field) and the AdjustedProcessorCycles value (in the FastCounter field) that can be used in the calculation of the corresponding core crystal clock value of the TSC packet.

When the next MTC after a pair of TSC/TMA is seen, the number of crystal clocks passed since the TSC packet can be calculated by subtracting the TMA.CTC value from the time indicated by the  $MTC_{Next}$  packet by

$CTC_{Delta}[15:0] = (CTC_{Next}[15:0] - TMA.CTC[15:0])$ , where  $CTC_{Next} = MTC_{Payload} \ll IA32\_RTIT\_CTL.MTCFreq$ .

The TMA.FastCounter field provides the fractional component of the TSC packet into the next crystal clock cycle.

CYC packets can provide further precision of an estimated timestamp value to many non-timing packets, by providing an indication of the time passed between other timing packets (MTCs or TSCs).

When enabled, CYC packets are sent preceding each CYC-eligible packet, and provide the number of processor core clock cycles that have passed since the last CYC packet. Thus between MTCs and TSCs, the accumulated CYC values can be used to estimate the adjusted\_processor\_cycles component of the timestamp value. The accumulated CPU cycles will have to be adjusted to account for the difference in frequency between the processor core clock and the P1 frequency. The necessary adjustment can be estimated using the core:bus ratio value given in the CBR packet, by multiplying the accumulated cycle count value by  $P1/CBR_{payload}$ .

Note that stand-alone TSC packets (that is, TSC packets that are not a part of a PSB+) are typically generated only when generation of other timing packets (MTCs and CYCs) has ceased for a period of time. Example scenarios include when Intel PT is re-enabled, or on wake after a sleep state. Thus any calculation of ART or cycle time leading up to a TSC packet will likely result in a discrepancy, which the TSC packet serves to correct.

A greater level of precision may be achieved by calculating the CPU clock frequency, see Section 35.8.3.4 below for a method to do so using Intel PT packets.

CYCs can be used to estimate time between TSCs even without MTCs, though this will likely result in a reduction in estimated TSC precision.

### 35.8.3.3 VMX TSC Manipulation

When software executes in non-Root operation, additional offset and scaling factors may be applied to the TSC value. These are optional, but may be enabled via VMCS controls on a per-VM basis. See Chapter 25, "VMX Non-Root Operation" for details on VMX TSC offsetting and TSC scaling.

Like the value returned by RDTSC, TSC packets will include these adjustments, but other timing packets (such as MTC, CYC, and CBR) are not impacted. In order to use the algorithm above to estimate the TSC value when TSC scaling is in use, it will be necessary for software to account for the scaling factor. See Section 35.5.2.6 for details.



### 35.8.3.4 Calculating Frequency with Intel PT

Because Intel PT can provide both wall-clock time and processor clock cycle time, it can be used to measure the processor core clock frequency. Either TSC or MTC packets can be used to track the wall-clock time. By using CYC packets to count the number of processor core cycles that pass in between a pair of wall-clock time packets, the ratio between processor core clock frequency and TSC frequency can be derived. If the P1 frequency is known, it can be applied to determine the CPU frequency. See Section 35.8.3.1 above for details on the relationship between TSC, MTC, and CYC.

## 15. Updates to Chapter 40, Volume 3D

Change bars show changes to Chapter 35 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3D: System Programming Guide, Part 4*.

-----  
Changes to this chapter: ENCLS, ENCLU and ENCLV instructions updated to indicate there is no CPUID feature flag associated with these instructions. EEXIT instruction updated.

## CHAPTER 40 SGX INSTRUCTION REFERENCES

This chapter describes the supervisor and user level instructions provided by Intel® Software Guard Extensions (Intel® SGX). In general, various functionality is encoded as leaf functions within the ENCLS (supervisor), ENCLU (user), and the ENCLV (virtualization operation) instruction mnemonics. Different leaf functions are encoded by specifying an input value in the EAX register of the respective instruction mnemonic.

### 40.1 INTEL® SGX INSTRUCTION SYNTAX AND OPERATION

ENCLS, ENCLU and ENCLV instruction mnemonics for all leaf functions are covered in this section.

For all instructions, the value of CS.D is ignored; addresses and operands are 64 bits in 64-bit mode and are otherwise 32 bits. Aside from EAX specifying the leaf number as input, each instruction leaf may require all or some subset of the RBX/RCX/RDX as input parameters. Some leaf functions may return data or status information in one or more of the general purpose registers.

#### 40.1.1 ENCLS Register Usage Summary

Table 40-1 summarizes the implicit register usage of supervisor mode enclave instructions.

**Table 40-1. Register Usage of Privileged Enclave Instruction Leaf Functions**

| Instr. Leaf | EAX       | RBX                | RCX              | RDX                 |
|-------------|-----------|--------------------|------------------|---------------------|
| ECREATE     | 00H (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) |                     |
| EADD        | 01H (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) |                     |
| EINIT       | 02H (In)  | SIGSTRUCT (In, EA) | SECS (In, EA)    | EINITTOKEN (In, EA) |
| EREMOVE     | 03H (In)  |                    | EPCPAGE (In, EA) |                     |
| EDBGGRD     | 04H (In)  | Result Data (Out)  | EPCPAGE (In, EA) |                     |
| EDBGWR      | 05H (In)  | Source Data (In)   | EPCPAGE (In, EA) |                     |
| EEXTEND     | 06H (In)  | SECS (In, EA)      | EPCPAGE (In, EA) |                     |
| ELDB        | 07H (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) | VERSION (In, EA)    |
| ELDU        | 08H (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) | VERSION (In, EA)    |
| EBLOCK      | 09H (In)  |                    | EPCPAGE (In, EA) |                     |
| EPA         | 0AH (In)  | PT_VA (In)         | EPCPAGE (In, EA) |                     |
| EWB         | 0BH (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) | VERSION (In, EA)    |
| ETRACK      | 0CH (In)  |                    | EPCPAGE (In, EA) |                     |
| EAUG        | 0DH (In)  | PAGEINFO (In, EA)  | EPCPAGE (In, EA) |                     |
| EMODPR      | 0EH (In)  | SECINFO (In, EA)   | EPCPAGE (In, EA) |                     |
| EMODT       | 0FH (In)  | SECINFO (In, EA)   | EPCPAGE (In, EA) |                     |
| ERDINFO     | 010H (In) | RDINFO (In, EA*)   | EPCPAGE (In, EA) |                     |
| ETRACKC     | 011H (In) |                    | EPCPAGE (In, EA) |                     |
| ELDBC       | 012H (In) | PAGEINFO (In, EA*) | EPCPAGE (In, EA) | VERSION (In, EA)    |
| ELDUC       | 013H (In) | PAGEINFO (In, EA*) | EPCPAGE (In, EA) | VERSION (In, EA)    |

EA: Effective Address

### 40.1.2 ENCLU Register Usage Summary

Table 40-2 summarizes the implicit register usage of user mode enclave instructions.

**Table 40-2. Register Usage of Unprivileged Enclave Instruction Leaf Functions**

| Instr. Leaf           | EAX            | RBX                 | RCX                 | RDY                 |
|-----------------------|----------------|---------------------|---------------------|---------------------|
| EReport               | 00H (In)       | TARGETINFO (In, EA) | REPORTDATA (In, EA) | OUTPUTDATA (In, EA) |
| EGetKey               | 01H (In)       | KEYREQUEST (In, EA) | KEY (In, EA)        |                     |
| EEnter                | 02H (In)       | TCS (In, EA)        | AEP (In, EA)        |                     |
|                       | RBX.CSSA (Out) |                     | Return (Out, EA)    |                     |
| EResume               | 03H (In)       | TCS (In, EA)        | AEP (In, EA)        |                     |
| EExit                 | 04H (In)       | Target (In, EA)     | Current AEP (Out)   |                     |
| EAccept               | 05H (In)       | SECINFO (In, EA)    | EPCPAGE (In, EA)    |                     |
| EModPE                | 06H (In)       | SECINFO (In, EA)    | EPCPAGE (In, EA)    |                     |
| EAcceptCopy           | 07H (In)       | SECINFO (In, EA)    | EPCPAGE (In, EA)    | EPCPAGE (In, EA)    |
| EA: Effective Address |                |                     |                     |                     |

### 40.1.3 ENCLV Register Usage Summary

Table 40-3 summarizes the implicit register usage of virtualization operation enclave instructions.

**Table 40-3. Register Usage of Virtualization Operation Enclave Instruction Leaf Functions**

| Instr. Leaf           | EAX      | RBX              | RCX              | RDY                    |
|-----------------------|----------|------------------|------------------|------------------------|
| EDECvirtChild         | 00H (In) | EPCPAGE (In, EA) | SECS (In, EA)    |                        |
| EINCvirtChild         | 01H (In) | EPCPAGE (In, EA) | SECS (In, EA)    |                        |
| ESETCONTEXT           | 02H (In) |                  | EPCPAGE (In, EA) | Context Value (In, EA) |
| EA: Effective Address |          |                  |                  |                        |

### 40.1.4 Information and Error Codes

Information and error codes are reported by various instruction leaf functions to show an abnormal termination of the instruction or provide information which may be useful to the developer. Table 40-4 shows the various codes and the instruction which generated the code. Details of the meaning of the code is provided in the individual instruction.

**Table 40-4. Error or Information Codes for Intel® SGX Instructions**

| Name                    | Value | Returned By  |
|-------------------------|-------|--|
| No Error                | 0     |  |
| SGX_INVALID_SIG_STRUCT  | 1     | EINIT  |
| SGX_INVALID_ATTRIBUTE   | 2     | EINIT, EGETKEY   |
| SGX_BLSTATE             | 3     | EBLOCK   |
| SGX_INVALID_MEASUREMENT | 4     | EINIT  |
| SGX_NOTBLOCKABLE        | 5     | EBLOCK   |
| SGX_PG_INVLD            | 6     | EBLOCK, ERDINFO, ETRACKC   |
| SGX_EPC_PAGE_CONFLICT   | 7     | EBLOCK, EMODPR, EMODT, ERDINFO, EDECvirtChild, EINCvirtChild, ELDBC, ELDUC, ESETCONTEXT, ETRACKC |

**Table 40-4. Error or Information Codes for Intel® SGX Instructions**

| Name                         | Value | Returned By                 |
|------------------------------|-------|-----------------------------|
| SGX_INVALID_SIGNATURE        | 8     | EINIT                       |
| SGX_MAC_COMPARE_FAIL         | 9     | ELDB, ELDU, ELDBC, ELDUC    |
| SGX_PAGE_NOT_BLOCKED         | 10    | EWB                         |
| SGX_NOT_TRACKED              | 11    | EWB, EACCEPT                |
| SGX_VA_SLOT_OCCUPIED         | 12    | EWB                         |
| SGX_CHILD_PRESENT            | 13    | EWB, EREMOVE                |
| SGX_ENCLAVE_ACT              | 14    | EREMOVE                     |
| SGX_ENTRYEPOCH_LOCKED        | 15    | EBLOCK                      |
| SGX_INVALID_EINITTOKEN       | 16    | EINIT                       |
| SGX_PREV_TRK_INCMPL          | 17    | ETRACK, ETRACKC             |
| SGX_PG_IS_SECS               | 18    | EBLOCK                      |
| SGX_PAGE_ATTRIBUTES_MISMATCH | 19    | EACCEPT, EACCEPTCOPY        |
| SGX_PAGE_NOT_MODIFIABLE      | 20    | EMODPR, EMODT               |
| SGX_PAGE_NOT_DEBUGGABLE      | 21    | EDBGRD, EDBGWR              |
| SGX_INVALID_COUNTER          | 25    | EDECVIRTCHILD, EINCVRTCHILD |
| SGX_PG_NONEPC                | 26    | ERDINFO                     |
| SGX_TRACK_NOT_REQUIRED       | 27    | ETRACKC                     |
| SGX_INVALID_CPUSVN           | 32    | EINIT, EGETKEY              |
| SGX_INVALID_ISVSVN           | 64    | EGETKEY                     |
| SGX_UNMASKED_EVENT           | 128   | EINIT                       |
| SGX_INVALID_KEYNAME          | 256   | EGETKEY                     |

### 40.1.5 Internal CREGs

The CREGs as shown in Table 5-4 are hardware specific registers used in this document to indicate values kept by the processor. These values are used while executing in enclave mode or while executing an Intel SGX instruction. These registers are not software visible and are implementation specific. The values in Table 40-5 appear at various places in the pseudo-code of this document. They are used to enhance understanding of the operations.

**Table 40-5. List of Internal CREG**

| Name                     | Size (Bits) | Scope |
|--------------------------|-------------|-------|
| CR_ENCLAVE_MODE          | 1           | LP    |
| CR_DBGOPTIN              | 1           | LP    |
| CR_TCS_LA                | 64          | LP    |
| CR_TCS_PA                | 64          | LP    |
| CR_ACTIVE_SECS           | 64          | LP    |
| CR_EL RANGE              | 128         | LP    |
| CR_SAVE_TF               | 1           | LP    |
| CR_SAVE_FS               | 64          | LP    |
| CR_GPR_PA                | 64          | LP    |
| CR_XSAVE_PAGE_n          | 64          | LP    |
| CR_SAVE_DR7              | 64          | LP    |
| CR_SAVE_PERF_GLOBAL_CTRL | 64          | LP    |

**Table 40-5. List of Internal CREG**

| Name                   | Size (Bits) | Scope   |
|------------------------|-------------|---------|
| CR_SAVE_DEBUGCTL       | 64          | LP      |
| CR_SAVE_PEBS_ENABLE    | 64          | LP      |
| CR_CPUSVN              | 128         | PACKAGE |
| CR_SGXOWNERPOCH        | 128         | PACKAGE |
| CR_SAVE_XCRO           | 64          | LP      |
| CR_SGX_ATTRIBUTES_MASK | 128         | LP      |
| CR_PAGING_VERSION      | 64          | PACKAGE |
| CR_VERSION_THRESHOLD   | 64          | PACKAGE |
| CR_NEXT_EID            | 64          | PACKAGE |
| CR_BASE_PK             | 128         | PACKAGE |
| CR_SEAL_FUSES          | 128         | PACKAGE |

### 40.1.6 Concurrent Operation Restrictions

Under certain conditions, Intel SGX disallows certain leaf functions from operating concurrently. Listed below are some examples of concurrency that are not allowed.

- For example, Intel SGX disallows the following leafs to concurrently operate on the same EPC page.
  - ECREATE, EADD, and EREMOVE are not allowed to operate on the same EPC page concurrently with themselves.
  - EADD, EEXTEND, and EINIT leaves are not allowed to operate on the same SECS concurrently.
- Intel SGX disallows the EREMOVE leaf from removing pages from an enclave that is in use.
- Intel SGX disallows entry (EENTER and ERESUME) to an enclave while a page from that enclave is being removed.

When disallowed operation is detected, a leaf function may do one of the following:

- Return an SGX\_EPC\_PAGE\_CONFLICT error code in RAX.
- Cause a #GP(0) exception.

To prevent such exceptions, software must serialize leaf functions or prevent these leaf functions from accessing the same EPC page.

#### 40.1.6.1 Concurrency Tables of Intel® SGX Instructions

The tables below detail the concurrent operation restrictions of all SGX leaf functions. For each leaf function, the table has a separate line for each of the EPC pages the leaf function accesses.

For each such EPC page, the base concurrency requirements are detailed as follows:

- **Exclusive Access** means that no other leaf function that requires either shared or exclusive access to the same EPC page may be executed concurrently. For example, EADD requires an exclusive access to the target page it accesses.
- **Shared Access** means that no other leaf function that requires an exclusive access to the same EPC page may be executed concurrently. Other leaf functions that require shared access may run concurrently. For example, EADD requires a shared access to the SECS page it accesses.
- **Concurrent Access** means that any other leaf function that requires any access to the same EPC page may be executed concurrently. For example, EGETKEY has no concurrency requirements for the KEYREQUEST page.

In addition to the base concurrency requirements, additional concurrency requirements are listed, which apply only to specific sets of leaf functions. For example, there are additional requirements that apply for EADD, EXTEND and EINIT. EADD and EEXTEND can't execute concurrently on the same SECS page.

The tables also detail the leaf function's behavior when a conflict happens, i.e., a concurrency requirement is not met. In this case, the leaf function may return an SGX\_EPC\_PAGE\_CONFLICT error code in RAX, or it may cause an exception. In addition, the tables detail those conflicts where a VM Exit may be triggered, and list the Exit Qualification code that is provided in such cases.

**Table 40-6. Base Concurrency Restrictions**

| Leaf          | Parameter  |                           | Base Concurrency Restrictions |                           |                                    |
|---------------|------------|---------------------------|-------------------------------|---------------------------|------------------------------------|
|               |            |                           | Access                        | On Conflict               | SGX_CONFLICT VM Exit Qualification |
| EACCEPT       | Target     | [DS:RCX]                  | Shared                        | #GP                       |                                    |
|               | SECINFO    | [DS:RBX]                  | Concurrent                    |                           |                                    |
| EACCEPTCOPY   | Target     | [DS:RCX]                  | Concurrent                    |                           |                                    |
|               | Source     | [DS:RDX]                  | Concurrent                    |                           |                                    |
|               | SECINFO    | [DS:RBX]                  | Concurrent                    |                           |                                    |
| EADD          | Target     | [DS:RCX]                  | Exclusive                     | #GP                       | EPC_PAGE_CONFLICT_EXCEPTION        |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Shared                        | #GP                       |                                    |
| EAUG          | Target     | [DS:RCX]                  | Exclusive                     | #GP                       | EPC_PAGE_CONFLICT_EXCEPTION        |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Shared                        | #GP                       |                                    |
| EBLOCK        | Target     | [DS:RCX]                  | Shared                        | SGX_EPC_PAGE<br>_CONFLICT |                                    |
| ECREATE       | SECS       | [DS:RCX]                  | Exclusive                     | #GP                       | EPC_PAGE_CONFLICT_EXCEPTION        |
| EDBGRD        | Target     | [DS:RCX]                  | Shared                        | #GP                       |                                    |
| EDBGWR        | Target     | [DS:RCX]                  | Shared                        | #GP                       |                                    |
| EDECVIRTCHILD | Target     | [DS:RBX]                  | Shared                        | SGX_EPC_PAGE<br>_CONFLICT |                                    |
|               | SECS       | [DS:RCX]                  | Concurrent                    |                           |                                    |
| EENTERTCS     | SECS       | [DS:RBX]                  | Shared                        | #GP                       |                                    |
| EEXIT         |            |                           | Concurrent                    |                           |                                    |
| EEXTEND       | Target     | [DS:RCX]                  | Shared                        | #GP                       |                                    |
|               | SECS       | [DS:RBX]                  | Concurrent                    |                           |                                    |
| EGETKEY       | KEYREQUEST | [DS:RBX]                  | Concurrent                    |                           |                                    |
|               | OUTPUTDATA | [DS:RCX]                  | Concurrent                    |                           |                                    |
| EINCVIRTCHILD | Target     | [DS:RBX]                  | Shared                        | SGX_EPC_PAGE<br>_CONFLICT |                                    |
|               | SECS       | [DS:RCX]                  | Concurrent                    |                           |                                    |
| EINIT         | SECS       | [DS:RCX]                  | Shared                        | #GP                       |                                    |
| ELDB/ELDU     | Target     | [DS:RCX]                  | Exclusive                     | #GP                       | EPC_PAGE_CONFLICT_EXCEPTION        |
|               | VA         | [DS:RDX]                  | Shared                        | #GP                       |                                    |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Shared                        | #GP                       |                                    |

**Table 40-6. Base Concurrency Restrictions**

| Leaf        | Parameter  |                           | Base Concurrency Restrictions |                       |                                    |
|-------------|------------|---------------------------|-------------------------------|-----------------------|------------------------------------|
|             |            |                           | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| EDLBC/ELDUC | Target     | [DS:RCX]                  | Exclusive                     | SGX_EPC_PAGE_CONFLICT | EPC_PAGE_CONFLICT_ERROR            |
|             | VA         | [DS:RDX]                  | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|             | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
| EMODPE      | Target     | [DS:RCX]                  | Concurrent                    |                       |                                    |
|             | SECINFO    | [DS:RBX]                  | Concurrent                    |                       |                                    |
| EMODPR      | Target     | [DS:RCX]                  | Shared                        | #GP                   |                                    |
| EMODT       | Target     | [DS:RCX]                  | Exclusive                     | SGX_EPC_PAGE_CONFLICT | EPC_PAGE_CONFLICT_ERROR            |
| EPA         | VA         | [DS:RCX]                  | Exclusive                     | #GP                   | EPC_PAGE_CONFLICT_EXCEPTION        |
| ERDINFO     | Target     | [DS:RCX]                  | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
| EREMOVE     | Target     | [DS:RCX]                  | Exclusive                     | #GP                   | EPC_PAGE_CONFLICT_EXCEPTION        |
| EREPORT     | TARGETINFO | [DS:RBX]                  | Concurrent                    |                       |                                    |
|             | REPORTDATA | [DS:RCX]                  | Concurrent                    |                       |                                    |
|             | OUTPUTDATA | [DS:RDX]                  | Concurrent                    |                       |                                    |
| ERESUME     | TCS        | [DS:RBX]                  | Shared                        | #GP                   |                                    |
| ESETCONTEXT | SECS       | [DS:RCX]                  | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
| ETRACK      | SECS       | [DS:RCX]                  | Shared                        | #GP                   |                                    |
| ETRACKC     | Target     | [DS:RCX]                  | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|             | SECS       | Implicit                  | Concurrent                    |                       |                                    |
| EWB         | Source     | [DS:RCX]                  | Exclusive                     | #GP                   | EPC_PAGE_CONFLICT_EXCEPTION        |
|             | VA         | [DS:RDX]                  | Shared                        | #GP                   |                                    |

**Table 40-7. Additional Concurrency Restrictions**

| Leaf        | Parameter |          | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------------|-----------|----------|---|-------------|--------------------------|-------------|---------------------|-------------|
|             |           |          | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|             |           |          | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EACCEPT     | Target    | [DS:RCX] | Exclusive                                       | #GP         | Concurrent               |             | Concurrent          |             |
|             | SECINFO   | [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
| EACCEPTCOPY | Target    | [DS:RCX] | Exclusive                                       | #GP         | Concurrent               |             | Concurrent          |             |
|             | Source    | [DS:RDX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | SECINFO   | [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |



Table 40-7. Additional Concurrency Restrictions

| Leaf          | Parameter  |                           | Additional Concurrency Restrictions             |                               |                          |             |                     |             |
|---------------|------------|---------------------------|---|-------------------------------|--------------------------|-------------|---------------------|-------------|
|               |            |                           | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |                               | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|               |            |                           | Access  | On Conflict                   | Access                   | On Conflict | Access              | On Conflict |
| EADD          | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Concurrent                                      |                               | Exclusive                | #GP         | Concurrent          |             |
| EAUG          | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EBLOCK        | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| ECREATE       | SECS       | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EDBGGRD       | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EDBGWR        | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EDECVRTCHILD  | Target     | [DS:RBX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EENTERTCS     | SECS       | [DS:RBX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EEXIT         |            |                           | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EEXTEND       | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RBX]                  | Concurrent                                      |                               | Exclusive                | #GP         | Concurrent          |             |
| EGETKEY       | KEYREQUEST | [DS:RBX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | OUTPUTDATA | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EINCVIRTCHILD | Target     | [DS:RBX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EINIT         | SECS       | [DS:RCX]                  | Concurrent                                      |                               | Exclusive                | #GP         | Concurrent          |             |
| ELDB/ELDU     | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | VA         | [DS:RDX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EDLBC/ELDUC   | Target     | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | VA         | [DS:RDX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
|               | SECS       | [DS:RBX]PAGEINFO.<br>SECS | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EMODPE        | Target     | [DS:RCX]                  | Exclusive                                       | #GP                           | Concurrent               |             | Concurrent          |             |
|               | SECINFO    | [DS:RBX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |
| EMODPR        | Target     | [DS:RCX]                  | Exclusive                                       | SGX_EPC_<br>PAGE_CON<br>FLICT | Concurrent               |             | Concurrent          |             |
| EMODT         | Target     | [DS:RCX]                  | Exclusive                                       | SGX_EPC_<br>PAGE_CON<br>FLICT | Concurrent               |             | Concurrent          |             |
| EPA           | VA         | [DS:RCX]                  | Concurrent                                      |                               | Concurrent               |             | Concurrent          |             |

**Table 40-7. Additional Concurrency Restrictions**

| Leaf        | Parameter  |          | Additional Concurrency Restrictions             |             |                          |             |                     |                                    |
|-------------|------------|----------|---|-------------|--------------------------|-------------|---------------------|------------------------------------|
|             |            |          | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |                                    |
|             |            |          | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict                        |
| ERDINFO     | Target     | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
| EREMOVE     | Target     | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
| EREPORT     | TARGETINFO | [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
|             | REPORTDATA | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
|             | OUTPUTDATA | [DS:RDX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
| ERESUME     | TCS        | [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
| ESETCONTEXT | SECS       | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
| ETRACK      | SECS       | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Exclusive           | SGX_EPC_PAGE_CONFLICT <sup>1</sup> |
| ETRACKC     | Target     | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
|             | SECS       | Implicit | Concurrent                                      |             | Concurrent               |             | Exclusive           | SGX_EPC_PAGE_CONFLICT <sup>1</sup> |
| EWB         | Source     | [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |
|             | VA         | [DS:RDX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                                    |

**NOTES:**

1. SGX\_CONFLICT VM Exit Qualification =TRACKING\_RESOURCE\_CONFLICT.

## 40.2 INTEL® SGX INSTRUCTION REFERENCE

## ENCLS—Execute an Enclave System Function of Specified Leaf Number

| Opcode/<br>Instruction | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|------------------------|-------|------------------------------|--------------------------|--|
| NP 0F 01 CF<br>ENCLS   | NP    | V/V                          | NA                       | This instruction is used to execute privileged Intel SGX leaf functions that are used for managing and debugging the enclaves. |

### Instruction Operand Encoding

| Op/En | Operand 1 | Operand 2 | Operand 3 | Implicit Register Operands |
|-------|-----------|-----------|-----------|----------------------------|
| NP    | NA        | NA        | NA        | See Section 40.3           |

### Description

The ENCLS instruction invokes the specified privileged Intel SGX leaf function for managing and debugging enclaves. Software specifies the leaf function by setting the appropriate value in the register EAX as input. The registers RBX, RCX, and RDX have leaf-specific purpose, and may act as input, as output, or may be unused. In 64-bit mode, the instruction ignores upper 32 bits of the RAX register.

The ENCLS instruction produces an invalid-opcode exception (#UD) if CR0.PE = 0 or RFLAGS.VM = 1, or if it is executed in system-management mode (SMM). Additionally, any attempt to execute the instruction when CPL > 0 results in #UD. The instruction produces a general-protection exception (#GP) if CR0.PG = 0 or if an attempt is made to invoke an undefined leaf function.

In VMX non-root operation, execution of ENCLS may cause a VM exit if the “enable ENCLS exiting” VM-execution control is 1. In this case, execution of individual leaf functions of ENCLS is governed by the ENCLS-exiting bitmap field in the VMCS. Each bit in that field corresponds to the index of an ENCLS leaf function (as provided in EAX).

Software in VMX root operation can thus intercept the invocation of various ENCLS leaf functions in VMX non-root operation by setting the “enable ENCLS exiting” VM-execution control and setting the corresponding bits in the ENCLS-exiting bitmap.

Addresses and operands are 32 bits outside 64-bit mode (IA32\_EFER.LMA = 0 || CS.L = 0) and are 64 bits in 64-bit mode (IA32\_EFER.LMA = 1 || CS.L = 1). CS.D value has no impact on address calculation. The DS segment is used to create linear addresses.

Segment override prefixes and address-size override prefixes are ignored, and is the REX prefix in 64-bit mode.

### Operation

IF TSX\_ACTIVE

THEN GOTO TSX\_ABORT\_PROCESSING; FI;

IF CR0.PE = 0 or RFLAGS.VM = 1 or in SMM or CPUID.SGX\_LEAF.0:EAX.SE1 = 0

THEN #UD; FI;

IF (CPL > 0)

THEN #UD; FI;

IF in VMX non-root operation and the “enable ENCLS exiting” VM-execution control is 1

THEN

IF EAX < 63 and ENCLS\_exiting\_bitmap[EAX] = 1 or EAX > 62 and ENCLS\_exiting\_bitmap[63] = 1

THEN VM exit;

FI;

FI;

IF IA32\_FEATURE\_CONTROL.LOCK = 0 or IA32\_FEATURE\_CONTROL.SGX\_ENABLE = 0

THEN #GP(0); FI;

IF EAX is invalid leaf number)

THEN #GP(0); FI;

IF CR0.PG = 0  
THEN #GP(0); FI;

(\* DS must not be an expanded down segment \*)  
IF not in 64-bit mode and DS.Type is expand-down data  
THEN #GP(0); FI;

Jump to leaf specific flow

### Flags Affected

See individual leaf functions

### Protected Mode Exceptions

#UD                    If any of the LOCK/OSIZE/REP/VEX prefix is used.  
                         If current privilege level is not 0.  
                         If CPUID.(EAX=12H,ECX=0):EAX.SGX1 [bit 0] = 0.  
                         If logical processor is in SMM.

#GP(0)                If IA32\_FEATURE\_CONTROL.LOCK = 0.  
                         If IA32\_FEATURE\_CONTROL.SGX\_ENABLE = 0.  
                         If input value in EAX encodes an unsupported leaf.  
                         If data segment expand down.  
                         If CR0.PG=0.

### Real-Address Mode Exceptions

#UD                    ENCLS is not recognized in real mode.

### Virtual-8086 Mode Exceptions

#UD                    ENCLS is not recognized in virtual-8086 mode.

### Compatibility Mode Exceptions

Same exceptions as in protected mode.

### 64-Bit Mode Exceptions

#UD                    If any of the LOCK/OSIZE/REP/VEX prefix is used.  
                         If current privilege level is not 0.  
                         If CPUID.(EAX=12H,ECX=0):EAX.SGX1 [bit 0] = 0.  
                         If logical processor is in SMM.

#GP(0)                If IA32\_FEATURE\_CONTROL.LOCK = 0.  
                         If IA32\_FEATURE\_CONTROL.SGX\_ENABLE = 0.  
                         If input value in EAX encodes an unsupported leaf.

## ENCLU—Execute an Enclave User Function of Specified Leaf Number

| Opcode/<br>Instruction | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|------------------------|-------|------------------------------|--------------------------|--|
| NP 0F 01 D7<br>ENCLU   | NP    | V/V                          | NA                       | This instruction is used to execute non-privileged Intel SGX leaf functions. |

### Instruction Operand Encoding

| Op/En | Operand 1 | Operand 2 | Operand 3 | Implicit Register Operands |
|-------|-----------|-----------|-----------|----------------------------|
| NP    | NA        | NA        | NA        | See Section 40.4           |

### Description

The ENCLU instruction invokes the specified non-privileged Intel SGX leaf functions. Software specifies the leaf function by setting the appropriate value in the register EAX as input. The registers RBX, RCX, and RDX have leaf-specific purpose, and may act as input, as output, or may be unused. In 64-bit mode, the instruction ignores upper 32 bits of the RAX register.

The ENCLU instruction produces an invalid-opcode exception (#UD) if CR0.PE = 0 or RFLAGS.VM = 1, or if it is executed in system-management mode (SMM). Additionally, any attempt to execute this instruction when CPL < 3 results in #UD. The instruction produces a general-protection exception (#GP) if either CR0.PG or CR0.NE is 0, or if an attempt is made to invoke an undefined leaf function. The ENCLU instruction produces a device not available exception (#NM) if CR0.TS = 1.

Addresses and operands are 32 bits outside 64-bit mode (IA32\_EFER.LMA = 0 or CS.L = 0) and are 64 bits in 64-bit mode (IA32\_EFER.LMA = 1 and CS.L = 1). CS.D value has no impact on address calculation. The DS segment is used to create linear addresses.

Segment override prefixes and address-size override prefixes are ignored, as is the REX prefix in 64-bit mode.

### Operation

```
IN_64BIT_MODE ← 0;
```

```
IF TSX_ACTIVE
```

```
    THEN GOTO TSX_ABORT_PROCESSING; FI;
```

```
IF CR0.PE = 0 or RFLAGS.VM = 1 or in SMM or CPUID.SGX_LEAF.0:EAX.SE1 = 0
```

```
    THEN #UD; FI;
```

```
IF CR0.TS = 1
```

```
    THEN #NM; FI;
```

```
IF CPL < 3
```

```
    THEN #UD; FI;
```

```
IF IA32_FEATURE_CONTROL.LOCK = 0 or IA32_FEATURE_CONTROL.SGX_ENABLE = 0
```

```
    THEN #GP(0); FI;
```

```
IF EAX is invalid leaf number
```

```
    THEN #GP(0); FI;
```

```
IF CR0.PG = 0 or CR0.NE = 0
```

```
    THEN #GP(0); FI;
```

```
IN_64BIT_MODE ← IA32_EFER.LMA AND CS.L ? 1 : 0;
```

```
(* Check not in 16-bit mode and DS is not a 16-bit segment *)
```

```
IF not in 64-bit mode and (CS.D = 0 or DS.B = 0)
```

THEN #GP(0); FI;

IF CR\_ENCLAVE\_MODE = 1 and (EAX = 2 or EAX = 3) (\* EENTER or ERESUME \*)  
 THEN #GP(0); FI;

IF CR\_ENCLAVE\_MODE = 0 and (EAX = 0 or EAX = 1 or EAX = 4 or EAX = 5 or EAX = 6 or EAX = 7)  
 (\* EREPORT, EGETKEY, EEXIT, EACCEPT, EMODPE, or EACCEPTCOPY \*)  
 THEN #GP(0); FI;

Jump to leaf specific flow

### Flags Affected

See individual leaf functions

### Protected Mode Exceptions

|        |  |
|--------|--|
| #UD    | If any of the LOCK/OSIZE/REP/VEX prefix is used.<br>If current privilege level is not 3.<br>If CPUID.(EAX=12H,ECX=0):EAX.SGX1 [bit 0] = 0.<br>If logical processor is in SMM.  |
| #GP(0) | If IA32_FEATURE_CONTROL.LOCK = 0.<br>If IA32_FEATURE_CONTROL.SGX_ENABLE = 0.<br>If input value in EAX encodes an unsupported leaf.<br>If input value in EAX encodes EENTER/ERESUME and ENCLAVE_MODE = 1.<br>If input value in EAX encodes EGETKEY/EREPORT/EEXIT/EACCEPT/EACCEPTCOPY/EMODPE and ENCLAVE_MODE = 0.<br>If operating in 16-bit mode.<br>If data segment is in 16-bit mode.<br>If CR0.PG = 0 or CR0.NE = 0. |
| #NM    | If CR0.TS = 1.   |

### Real-Address Mode Exceptions

|     |                                       |
|-----|---------------------------------------|
| #UD | ENCLS is not recognized in real mode. |
|-----|---------------------------------------|

### Virtual-8086 Mode Exceptions

|     |   |
|-----|---|
| #UD | ENCLS is not recognized in virtual-8086 mode. |
|-----|---|

### Compatibility Mode Exceptions

Same exceptions as in protected mode.

**64-Bit Mode Exceptions**

|        |  |
|--------|--|
| #UD    | If any of the LOCK/OSIZE/REP/VEX prefix is used.<br>If current privilege level is not 3.<br>If CPUID.(EAX=12H,ECX=0):EAX.SGX1 [bit 0] = 0.<br>If logical processor is in SMM.  |
| #GP(0) | If IA32_FEATURE_CONTROL.LOCK = 0.<br>If IA32_FEATURE_CONTROL.SGX_ENABLE = 0.<br>If input value in EAX encodes an unsupported leaf.<br>If input value in EAX encodes EENTER/ERESUME and ENCLAVE_MODE = 1.<br>If input value in EAX encodes EGETKEY/EREPORT/EEXIT/EACCEPT/EACCEPTCOPY/EMODPE and ENCLAVE_MODE = 0.<br>If CR0.NE = 0. |
| #NM    | If CR0.TS = 1.   |

## ENCLV—Execute an Enclave VMM Function of Specified Leaf Number

| Opcode/<br>Instruction | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|------------------------|-------|------------------------------|--------------------------|---|
| NP OF 01 C0<br>ENCLV   | NP    | V/V                          | NA                       | This instruction is used to execute privileged SGX leaf functions that are reserved for VMM use. They are used for managing the enclaves. |

### Instruction Operand Encoding

| Op/En | Operand 1 | Operand 2 | Operand 3 | Implicit Register Operands |
|-------|-----------|-----------|-----------|----------------------------|
| NP    | NA        | NA        | NA        | See Section 40.3           |

### Description

The ENCLV instruction invokes the virtualization SGX leaf functions for managing enclaves in a virtualized environment. Software specifies the leaf function by setting the appropriate value in the register EAX as input. The registers RBX, RCX, and RDX have leaf-specific purpose, and may act as input, as output, or may be unused. In non 64-bit mode, the instruction ignores upper 32 bits of the RAX register.

The ENCLV instruction produces an invalid-opcode exception (#UD) if CR0.PE = 0 or RFLAGS.VM = 1, if it is executed in system-management mode (SMM), or not in VMX operation. Additionally, any attempt to execute the instruction when CPL > 0 results in #UD. The instruction produces a general-protection exception (#GP) if CR0.PG = 0 or if an attempt is made to invoke an undefined leaf function.

Software in VMX root mode of operation can enable execution of the ENCLV instruction in VMX non-root mode by setting enable ENCLV execution control in the VMCS. If enable ENCLV execution control in the VMCS is clear, execution of the ENCLV instruction in VMX non-root mode results in #UD.

When execution of ENCLV instruction in VMX non-root mode is enabled, software in VMX root operation can intercept the invocation of various ENCLS leaf functions in VMX non-root operation by setting the corresponding bits in the ENCLV-exiting bitmap.

Addresses and operands are 32 bits in 32-bit mode (IA32\_EFER.LMA == 0 || CS.L == 0) and are 64 bits in 64-bit mode (IA32\_EFER.LMA == 1 && CS.L == 1). CS.D value has no impact on address calculation.

Segment override prefixes and address-size override prefixes are ignored, as is the REX prefix in 64-bit mode.

### Operation

```

IF TSX_ACTIVE
    THEN GOTO TSX_ABORT_PROCESSING; FI;

IF CR0.PE = 0 or RFLAGS.VM = 1 or in SMM or CPUID.SGX_LEAF.0:EAX.OSS = 0
    THEN #UD; FI;

IF in VMX non-root operation and IA_32_EFER.LMA = 1 and CS.L = 1
    THEN #UD; FI;

IF (CPL > 0)
    THEN #UD; FI;

IF in VMX non-root operation
    IF "enable ENCLV exiting" VM-execution control is 1
        THEN
            IF EAX < 63 and ENCLV_exiting_bitmap[EAX] = 1 or EAX > 62 and ENCLV_exiting_bitmap[63] = 1
                THEN VM exit;
            FI;
        ELSE
            #UD; FI;
    
```



FI;

IF IA32\_FEATURE\_CONTROL.LOCK = 0 or IA32\_FEATURE\_CONTROL.SGX\_ENABLE = 0  
THEN #GP(0); FI;

IF EAX is invalid leaf number)  
THEN #GP(0); FI;

IF CR0.PG = 0  
THEN #GP(0); FI;

(\* DS must not be an expanded down segment \*)  
IF not in 64-bit mode and DS.Type is expand-down data  
THEN #GP(0); FI;

Jump to leaf specific flow

### Flags Affected

See individual leaf functions.

### Protected Mode Exceptions

|        |  |
|--------|--|
| #UD    | If any of the LOCK/OSIZE/REP/VEX prefix is used.<br>If current privilege level is not 0.<br>If CPUID.(EAX=12H,ECX=0):EAX.OSS [bit 5] = 0.<br>If logical processor is in SMM.       |
| #GP(0) | If IA32_FEATURE_CONTROL.LOCK = 0.<br>If IA32_FEATURE_CONTROL.SGX_ENABLE = 0.<br>If input value in EAX encodes an unsupported leaf.<br>If data segment expand down.<br>If CR0.PG=0. |

### Real-Address Mode Exceptions

|     |                                       |
|-----|---------------------------------------|
| #UD | ENCLV is not recognized in real mode. |
|-----|---------------------------------------|

### Virtual-8086 Mode Exceptions

|     |   |
|-----|---|
| #UD | ENCLV is not recognized in virtual-8086 mode. |
|-----|---|

### Compatibility Mode Exceptions

Same exceptions as in protected mode.

### 64-Bit Mode Exceptions

|        |  |
|--------|--|
| #UD    | If any of the LOCK/OSIZE/REP/VEX prefix is used.<br>If current privilege level is not 0.<br>If CPUID.(EAX=12H,ECX=0):EAX.OSS [bit 5] = 0.<br>If logical processor is in SMM. |
| #GP(0) | If IA32_FEATURE_CONTROL.LOCK = 0.<br>If IA32_FEATURE_CONTROL.SGX_ENABLE = 0.<br>If input value in EAX encodes an unsupported leaf.   |

## 40.3 INTEL® SGX SYSTEM LEAF FUNCTION REFERENCE

Leaf functions available with the ENCLS instruction mnemonic are covered in this section. In general, each instruction leaf requires EAX to specify the leaf function index and/or additional implicit registers specifying leaf-specific input parameters. An instruction operand encoding table provides details of each implicit register usage and associated input/output semantics.

In many cases, an input parameter specifies an effective address associated with a memory object inside or outside the EPC, the memory addressing semantics of these memory objects are also summarized in a separate table.

## EADD—Add a Page to an Uninitialized Enclave

| Opcode/<br>Instruction   | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|--------------------------|-------|------------------------------|--------------------------|---|
| EAX = 01H<br>ENCLS[EADD] | IR    | V/V                          | SGX1                     | This leaf function adds a page to an uninitialized enclave. |

### Instruction Operand Encoding

| Op/En | EAX       | RBX                        | RCX                                      |
|-------|-----------|----------------------------|--|
| IR    | EADD (In) | Address of a PAGEINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function copies a source page from non-enclave memory into the EPC, associates the EPC page with an SECS page residing in the EPC, and stores the linear address and security attributes in EPCM. As part of the association, the enclave offset and the security attributes are measured and extended into the SECS.MRENCLAVE. This instruction can only be executed when current privilege level is 0.

RBX contains the effective address of a PAGEINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of EADD leaf function.

### EADD Memory Parameter Semantics

| PAGEINFO                             | PAGEINFO.SECS                          | PAGEINFO.SRCPGE                      | PAGEINFO.SECINFO                     | EPCPAGE                           |
|--------------------------------------|--|--------------------------------------|--------------------------------------|-----------------------------------|
| Read access permitted by Non Enclave | Read/Write access permitted by Enclave | Read access permitted by Non Enclave | Read access permitted by Non Enclave | Write access permitted by Enclave |

The instruction faults if any of the following:

### EADD Faulting Conditions

|   |  |
|---|--|
| The operands are not properly aligned.    | Unsupported security attributes are set.   |
| Refers to an invalid SECS.                | Reference is made to an SECS that is locked by another thread.   |
| The EPC page is locked by another thread. | RCX does not contain an effective address of an EPC page.  |
| The EPC page is already valid.            | If security attributes specifies a TCS and the source page specifies unsupported TCS values or fields. |
| The SECS has been initialized.            | The specified enclave offset is outside of the enclave address space.                                  |

### Concurrency Restrictions

**Table 40-8. Base Concurrency Restrictions of EADD**

| Leaf | Parameter                  | Base Concurrency Restrictions |             |                                    |
|------|----------------------------|-------------------------------|-------------|------------------------------------|
|      |                            | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EADD | Target [DS:RCX]            | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |
|      | SECS [DS:RBX]PAGEINFO.SECS | Shared                        | #GP         |                                    |

**Table 40-9. Additional Concurrency Restrictions of EADD**

| Leaf | Parameter                   | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|------|-----------------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|      |                             | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|      |                             | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EADD | Target [DS:RCX]             | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|      | SECS [DS:RBX]PAGE-INFO.SECS | Concurrent                                      |             | Exclusive                | #GP         | Concurrent          |             |

**Operation**

**Temp Variables in EADD Operational Flow**

| Name              | Type              | Size (bits) | Description   |
|-------------------|-------------------|-------------|---|
| TMP_SRCPGE        | Effective Address | 32/64       | Effective address of the source page.   |
| TMP_SECS          | Effective Address | 32/64       | Effective address of the SECS destination page.   |
| TMP_SECINFO       | Effective Address | 32/64       | Effective address of an SECINFO structure which contains security attributes of the page to be added. |
| SCRATCH_SECINFO   | SECINFO           | 512         | Scratch storage for holding the contents of DS:TMP_SECINFO.   |
| TMP_LINADDR       | Unsigned Integer  | 64          | Holds the linear address to be stored in the EPCM and used to calculate TMP_ENCLAVEOFFSET.            |
| TMP_ENCLAVEOFFSET | Enclave Offset    | 64          | The page displacement from the enclave base address.  |
| TMPUPDATEFIELD    | SHA256 Buffer     | 512         | Buffer used to hold data being added to TMP_SECS.MRENCLAVE.   |

IF (DS:RBX is not 32Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

TMP\_SRCPGE ← DS:RBX.SRCPGE;  
TMP\_SECS ← DS:RBX.SECS;  
TMP\_SECINFO ← DS:RBX.SECINFO;  
TMP\_LINADDR ← DS:RBX.LINADDR;

IF (DS:TMP\_SRCPGE is not 4KByte aligned or DS:TMP\_SECS is not 4KByte aligned or DS:TMP\_SECINFO is not 64Byte aligned or TMP\_LINADDR is not 4KByte aligned)  
THEN #GP(0); FI;

IF (DS:TMP\_SECS does not resolve within an EPC)  
THEN #PF(DS:TMP\_SECS); FI;

SCRATCH\_SECINFO ← DS:TMP\_SECINFO;

(\* Check for mis-configured SECINFO flags\*)  
IF (SCRATCH\_SECINFO reserved fields are not zero or

```

!(SCRATCH_SECINFO.FLAGS.PT is PT_REG or SCRATCH_SECINFO.FLAGS.PT is PT_TCS)
THEN #GP(0); FI;

(* Check the EPC page for concurrency *)
IF (SECS is not available for EADD)
THEN
    IF (<<VMX non-root operation>> AND <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
    THEN
        VMCS.Exit_reason ← SGX_CONFLICT;
        VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_EXCEPTION;
        VMCS.Exit_qualification.error ← 0;
        VMCS.Guest-physical_address ← << translation of DS:RCX produced by paging >>;
        VMCS.Guest-linear_address ← DS:RCX;
        Deliver VMEXIT;
    ELSE
        #GP(0);
    FI;
FI;

IF (EPCM(DS:RCX).VALID ≠ 0)
THEN #PF(DS:RCX); FI;

(* Check the SECS for concurrency *)
IF (SECS is not available for EADD)
THEN #GP(0); FI;

IF (EPCM(DS:TMP_SECS).VALID = 0 or EPCM(DS:TMP_SECS).PT ≠ PT_SECS)
THEN #PF(DS:TMP_SECS); FI;

(* Copy 4KBytes from source page to EPC page*)
DS:RCX[32767:0] ← DS:TMP_SRCPAGE[32767:0];

CASE (SCRATCH_SECINFO.FLAGS.PT)
{
    PT_TCS:
        IF (DS:RCX.RESERVED ≠ 0) #GP(0); FI;
        IF ( (DS:TMP_SECS.ATTRIBUTES.MODE64BIT = 0) and
            ((DS:TCS.FSLIMIT & 0FFFH ≠ 0FFFH) or (DS:TCS.GSLIMIT & 0FFFH ≠ 0FFFH)) ) #GP(0); FI;
        BREAK;
    PT_REG:
        IF (SCRATCH_SECINFO.FLAGS.W = 1 and SCRATCH_SECINFO.FLAGS.R = 0) #GP(0); FI;
        BREAK;
ESAC;

(* Check the enclave offset is within the enclave linear address space *)
IF (TMP_LINADDR < DS:TMP_SECS.BASEADDR or TMP_LINADDR ≥ DS:TMP_SECS.BASEADDR + DS:TMP_SECS.SIZE)
THEN #GP(0); FI;

(* Check concurrency of measurement resource*)
IF (Measurement being updated)
THEN #GP(0); FI;

(* Check if the enclave to which the page will be added is already in Initialized state *)
IF (DS:TMP_SECS already initialized)

```

```
THEN #GP(0); FI;
```

(\* For TCS pages, force EPCM.rwx bits to 0 and no debug access \*)

```
IF (SCRATCH_SECINFO.FLAGS.PT = PT_TCS)
```

```
THEN
```

```
    SCRATCH_SECINFO.FLAGS.R ← 0;
```

```
    SCRATCH_SECINFO.FLAGS.W ← 0;
```

```
    SCRATCH_SECINFO.FLAGS.X ← 0;
```

```
    (DS:RCX).FLAGS.DBGOPTIN ← 0; // force TCS.FLAGS.DBGOPTIN off
```

```
    DS:RCX.CSSA ← 0;
```

```
    DS:RCX.AEP ← 0;
```

```
    DS:RCX.STATE ← 0;
```

```
FI;
```

(\* Add enclave offset and security attributes to MRENCLAVE \*)

```
TMP_ENCLAVEOFFSET ← TMP_LINADDR - DS:TMP_SECS.BASEADDR;
```

```
TMPUPDATEFIELD[63:0] ← 0000000044444145H; // "EADD"
```

```
TMPUPDATEFIELD[127:64] ← TMP_ENCLAVEOFFSET;
```

```
TMPUPDATEFIELD[511:128] ← SCRATCH_SECINFO[375:0]; // 48 bytes
```

```
DS:TMP_SECS.MRENCLAVE ← SHA256UPDATE(DS:TMP_SECS.MRENCLAVE, TMPUPDATEFIELD)
```

```
INC enclave's MRENCLAVE update counter;
```

(\* Add enclave offset and security attributes to MRENCLAVE \*)

```
EPCM(DS:RCX).R ← SCRATCH_SECINFO.FLAGS.R;
```

```
EPCM(DS:RCX).W ← SCRATCH_SECINFO.FLAGS.W;
```

```
EPCM(DS:RCX).X ← SCRATCH_SECINFO.FLAGS.X;
```

```
EPCM(DS:RCX).PT ← SCRATCH_SECINFO.FLAGS.PT;
```

```
EPCM(DS:RCX).ENCLAVEADDRESS ← TMP_LINADDR;
```

(\* associate the EPCPAGE with the SECS by storing the SECS identifier of DS:TMP\_SECS \*)

```
Update EPCM(DS:RCX) SECS identifier to reference DS:TMP_SECS identifier;
```

(\* Set EPCM entry fields \*)

```
EPCM(DS:RCX).BLOCKED ← 0;
```

```
EPCM(DS:RCX).PENDING ← 0;
```

```
EPCM(DS:RCX).MODIFIED ← 0;
```

```
EPCM(DS:RCX).VALID ← 1;
```

### Flags Affected

None

### Protected Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand effective address is outside the DS segment limit.</li> <li>If a memory operand is not properly aligned.</li> <li>If an enclave memory operand is outside of the EPC.</li> <li>If an enclave memory operand is the wrong type.</li> <li>If a memory operand is locked.</li> <li>If the enclave is initialized.</li> <li>If the enclave's MRENCLAVE is locked.</li> <li>If the TCS page reserved bits are set.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If the EPC page is valid.</li> </ul>   |

**64-Bit Mode Exceptions**

|                 |   |
|-----------------|---|
| #GP(0)          | If a memory operand is non-canonical form.<br>If a memory operand is not properly aligned.<br>If an enclave memory operand is outside of the EPC.<br>If an enclave memory operand is the wrong type.<br>If a memory operand is locked.<br>If the enclave is initialized.<br>If the enclave's MRENCLAVE is locked.<br>If the TCS page reserved bits are set. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If the EPC page is valid.   |

## EAUG—Add a Page to an Initialized Enclave

| Opcode/<br>Instruction   | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|--------------------------|-------|------------------------------|--------------------------|---|
| EAX = 0DH<br>ENCLS[EAUG] | IR    | V/V                          | SGX2                     | This leaf function adds a page to an initialized enclave. |

### Instruction Operand Encoding

| Op/En | EAX       | RBX                        | RCX                                      |
|-------|-----------|----------------------------|--|
| IR    | EAUG (In) | Address of a SECFINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function zeroes a page of EPC memory, associates the EPC page with an SECS page residing in the EPC, and stores the linear address and security attributes in the EPCM. As part of the association, the security attributes are configured to prevent access to the EPC page until a corresponding invocation of the EACCEPT leaf or EACCEPT-COPY leaf confirms the addition of the new page into the enclave. This instruction can only be executed when current privilege level is 0.

RBX contains the effective address of a PAGEINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of the EAUG leaf function.

### EAUG Memory Parameter Semantics

| PAGEINFO                             | PAGEINFO.SECS                          | PAGEINFO.SRCPGE | PAGEINFO.SECINFO                     | EPCPAGE                           |
|--------------------------------------|--|-----------------|--------------------------------------|-----------------------------------|
| Read access permitted by Non Enclave | Read/Write access permitted by Enclave | Must be zero    | Read access permitted by Non Enclave | Write access permitted by Enclave |

The instruction faults if any of the following:

### EAUG Faulting Conditions

|   |   |
|---|---|
| The operands are not properly aligned.    | Unsupported security attributes are set.                              |
| Refers to an invalid SECS.                | Reference is made to an SECS that is locked by another thread.        |
| The EPC page is locked by another thread. | RCX does not contain an effective address of an EPC page.             |
| The EPC page is already valid.            | The specified enclave offset is outside of the enclave address space. |
| The SECS has been initialized.            |   |

### Concurrency Restrictions

**Table 40-10. Base Concurrency Restrictions of EAUG**

| Leaf | Parameter                  | Base Concurrency Restrictions |             |                                    |
|------|----------------------------|-------------------------------|-------------|------------------------------------|
|      |                            | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EAUG | Target [DS:RCX]            | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |
|      | SECS [DS:RBX]PAGEINFO.SECS | Shared                        | #GP         |                                    |



**Table 40-11. Additional Concurrency Restrictions of EAUG**

| Leaf | Parameter                   | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|------|-----------------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|      |                             | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|      |                             | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EAUG | Target [DS:RCX]             | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|      | SECS [DS:RBX]PAGE-INFO.SECS | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation****Temp Variables in EAUG Operational Flow**

| Name            | Type              | Size (bits) | Description   |
|-----------------|-------------------|-------------|---|
| TMP_SECS        | Effective Address | 32/64       | Effective address of the SECS destination page.   |
| TMP_SECINFO     | Effective Address | 32/64       | Effective address of an SECINFO structure which contains security attributes of the page to be added. |
| SCRATCH_SECINFO | SECINFO           | 512         | Scratch storage for holding the contents of DS:TMP_SECINFO.   |
| TMP_LINADDR     | Unsigned Integer  | 64          | Holds the linear address to be stored in the EPCM and used to calculate TMP_ENCLAVEOFFSET.            |

IF (DS:RBX is not 32Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

TMP\_SECS ← DS:RBX.SECS;  
TMP\_LINADDR ← DS:RBX.LINADDR;

IF ( DS:TMP\_SECS is not 4KByte aligned or TMP\_LINADDR is not 4KByte aligned )  
THEN #GP(0); FI;

IF ( ( DS:RBX.SRCPAGE is not 0 ) or ( DS:RBX.SECINFO is not 0 ) )  
THEN #GP(0); FI;

IF (DS:TMP\_SECS does not resolve within an EPC)  
THEN #PF(DS:TMP\_SECS); FI;

(\* Check the EPC page for concurrency \*)  
IF (EPC page in use)  
THEN #GP(0); FI;

IF (EPCM(DS:RCX).VALID ≠ 0)  
THEN #PF(DS:RCX); FI;

(\* Check the SECS for concurrency \*)

## SGX INSTRUCTION REFERENCES

IF (SECS is not available for EAUG)  
THEN #GP(0); FI;

IF (EPCM(DS:TMP\_SECS).VALID = 0 or EPCM(DS:TMP\_SECS).PT ≠ PT\_SECS)  
THEN #PF(DS:TMP\_SECS); FI;

(\* Check if the enclave to which the page will be added is in the Initialized state \*)  
IF (DS:TMP\_SECS is not initialized)  
THEN #GP(0); FI;

(\* Check the enclave offset is within the enclave linear address space \*)  
IF ( (TMP\_LINADDR < DS:TMP\_SECS.BASEADDR) or (TMP\_LINADDR ≥ DS:TMP\_SECS.BASEADDR + DS:TMP\_SECS.SIZE) )  
THEN #GP(0); FI;

(\* Clear the content of EPC page\*)  
DS:RCX[32767:0] ← 0;

(\* Set EPCM security attributes \*)  
EPCM(DS:RCX).R ← 1;  
EPCM(DS:RCX).W ← 1;  
EPCM(DS:RCX).X ← 0;  
EPCM(DS:RCX).PT ← PT\_REG;  
EPCM(DS:RCX).ENCLAVEADDRESS ← TMP\_LINADDR;  
EPCM(DS:RCX).BLOCKED ← 0;  
EPCM(DS:RCX).PENDING ← 1;  
EPCM(DS:RCX).MODIFIED ← 0;  
EPCM(DS:RCX).PR ← 0;

(\* associate the EPCPAGE with the SECS by storing the SECS identifier of DS:TMP\_SECS \*)  
Update EPCM(DS:RCX) SECS identifier to reference DS:TMP\_SECS identifier;

(\* Set EPCM valid fields \*)  
EPCM(DS:RCX).VALID ← 1;

### Flags Affected

None

### Protected Mode Exceptions

#GP(0)            If a memory operand effective address is outside the DS segment limit.  
                  If a memory operand is not properly aligned.  
                  If a memory operand is locked.  
                  If the enclave is not initialized.

#PF(error code)    If a page fault occurs in accessing memory operands.

### 64-Bit Mode Exceptions

#GP(0)            If a memory operand is non-canonical form.  
                  If a memory operand is not properly aligned.  
                  If a memory operand is locked.  
                  If the enclave is not initialized.

#PF(error code)    If a page fault occurs in accessing memory operands.

## EBLOCK—Mark a page in EPC as Blocked

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|----------------------------|-------|------------------------------|--------------------------|--|
| EAX = 09H<br>ENCLS[EBLOCK] | IR    | V/V                          | SGX1                     | This leaf function marks a page in the EPC as blocked. |

### Instruction Operand Encoding

| Op/En | EAX         |                         | RCX                                    |
|-------|-------------|-------------------------|--|
| IR    | EBLOCK (In) | Return error code (Out) | Effective address of the EPC page (In) |

### Description

This leaf function causes an EPC page to be marked as BLOCKED. This instruction can only be executed when current privilege level is 0.

The content of RCX is an effective address of an EPC page. The DS segment is used to create linear address. Segment override is not supported.

An error code is returned in RAX.

The table below provides additional information on the memory parameter of EBLOCK leaf function.

### EBLOCK Memory Parameter Semantics

|  |
|--|
| EPCPAGE                                |
| Read/Write access permitted by Enclave |

The error codes are:

**Table 40-12. EBLOCK Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EBLOCK successful.  |
| SGX_BLKSTATE                | Page already blocked. This value is used to indicate to a VMM that the page was already in BLOCKED state as a result of EBLOCK and thus will need to be restored to this state when it is eventually reloaded (using ELDB). |
| SGX_ENTRYEPOCH_LOCKED       | SECS locked for Entry Epoch update. This value indicates that an ETRACK is currently executing on the SECS. The EBLOCK should be reattempted.   |
| SGX_NOTBLOCKABLE            | Page type is not one which can be blocked.  |
| SGX_PG_INVLD                | Page is not valid and cannot be blocked.  |
| SGX_EPC_PAGE_CONFLICT       | Page is being written by EADD, EAUG, ECREATE, ELDU/B, EMODT, or EWB.  |

### Concurrency Restrictions

**Table 40-13. Base Concurrency Restrictions of EBLOCK**

| Leaf   | Parameter       | Base Concurrency Restrictions |                       |                                    |
|--------|-----------------|-------------------------------|-----------------------|------------------------------------|
|        |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| EBLOCK | Target [DS:RCX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |

**Table 40-14. Additional Concurrency Restrictions of EBLOCK**

| Leaf   | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|        |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EBLOCK | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EBLOCK Operational Flow**

| Name         | Type    | Size (Bits) | Description              |
|--------------|---------|-------------|--------------------------|
| TMP_BLKSTATE | Integer | 64          | Page is already blocked. |

IF (DS:RCX is not 4KByte Aligned)  
 THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

RFLAGS.ZF,CF,PF,AF,OF,SF ← 0;  
 RAX ← 0;

(\* Check the EPC page for concurrency\*)

IF (EPC page in use)  
 THEN  
     RFLAGS.ZF ← 1;  
     RAX ← SGX\_EPC\_PAGE\_CONFLICT;  
     GOTO DONE;

FI;

IF (EPCM(DS:RCX).VALID = 0)  
 THEN  
     RFLAGS.ZF ← 1;  
     RAX ← SGX\_PG\_INVLD;  
     GOTO DONE;

FI;

IF ( (EPCM(DS:RCX).PT ≠ PT\_REG) and (EPCM(DS:RCX).PT ≠ PT\_TCS) and (EPCM(DS:RCX).PT ≠ PT\_TRIM) )  
 THEN  
     RFLAGS.CF ← 1;  
     IF (EPCM(DS:RCX).PT = PT\_SECS)  
         THEN RAX ← SGX\_PG\_IS\_SECS;  
         ELSE RAX ← SGX\_NOTBLOCKABLE;  
     FI;  
     GOTO DONE;

FI;

(\* Check if the page is already blocked and report blocked state \*)  
 TMP\_BLKSTATE ← EPCM(DS:RCX).BLOCKED;

```

(* at this point, the page must be valid and PT_TCS or PT_REG or PT_TRIM*)
IF (TMP_BLKSTATE = 1)
  THEN
    RFLAGS.CF ← 1;
    RAX ← SGX_BLKSTATE;
  ELSE
    EPCM(DS:RCX).BLOCKED ← 1
FI;
DONE:

```

### Flags Affected

Sets ZF if SECS is in use or invalid, otherwise cleared. Sets CF if page is BLOCKED or not blockable, otherwise cleared. Clears PF, AF, OF, SF.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the DS segment limit.<br>If a memory operand is not properly aligned.<br>If the specified EPC resource is in use. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.  |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand is non-canonical form.<br>If a memory operand is not properly aligned.<br>If the specified EPC resource is in use. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.  |

## ECREATE—Create an SECS page in the Enclave Page Cache

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 00H<br>ENCLS[ECREATE] | IR    | V/V                          | SGX1                     | This leaf function begins an enclave build by creating an SECS page in EPC. |

### Instruction Operand Encoding

| Op/En | EAX          | RBX                        | RCX                                       |
|-------|--------------|----------------------------|---|
| IR    | ECREATE (In) | Address of a PAGEINFO (In) | Address of the destination SECS page (In) |

#### Description

ENCLS[ECREATE] is the first instruction executed in the enclave build process. ECREATE copies an SECS structure outside the EPC into an SECS page inside the EPC. The internal structure of SECS is not accessible to software.

ECREATE will set up fields in the protected SECS and mark the page as valid inside the EPC. ECREATE initializes or checks unused fields.

Software sets the following fields in the source structure: SECS:BASEADDR, SECS:SIZE in bytes, ATTRIBUTES, CONFIGID and CONFIGSVN. SECS:BASEADDR must be naturally aligned on an SECS.SIZE boundary. SECS.SIZE must be at least 2 pages (8192).

The source operand RBX contains an effective address of a PAGEINFO structure. PAGEINFO contains an effective address of a source SECS and an effective address of an SECINFO. The SECS field in PAGEINFO is not used.

The RCX register is the effective address of the destination SECS. It is an address of an empty slot in the EPC. The SECS structure must be page aligned. SECINFO flags must specify the page as an SECS page.

### ECREATE Memory Parameter Semantics

| PAGEINFO                             | PAGEINFO.SRCPGE                      | PAGEINFO.SECINFO                     | EPCPAGE                           |
|--------------------------------------|--------------------------------------|--------------------------------------|-----------------------------------|
| Read access permitted by Non Enclave | Read access permitted by Non Enclave | Read access permitted by Non Enclave | Write access permitted by Enclave |

ECREATE will fault if the SECS target page is in use; already valid; outside the EPC. It will also fault if addresses are not aligned; unused PAGEINFO fields are not zero.

If the amount of space needed to store the SSA frame is greater than the amount specified in SECS.SSAFRAME-SIZE, a #GP(0) results. The amount of space needed for an SSA frame is computed based on DS:TMP\_SECS.ATTRIBUTES.XFRM size. Details of computing the size can be found Section 41.7.

#### Concurrency Restrictions

**Table 40-15. Base Concurrency Restrictions of ECREATE**

| Leaf    | Parameter     | Base Concurrency Restrictions |             |                                    |
|---------|---------------|-------------------------------|-------------|------------------------------------|
|         |               | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| ECREATE | SECS [DS:RCX] | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |

Table 40-16. Additional Concurrency Restrictions of ECREATE

| Leaf    | Parameter     | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|---------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |               | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |               | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| ECREATE | SECS [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

## Operation

### Temp Variables in ECREATE Operational Flow

| Name           | Type              | Size (Bits) | Description  |
|----------------|-------------------|-------------|--|
| TMP_SRCPGE     | Effective Address | 32/64       | Effective address of the SECS source page.   |
| TMP_SECS       | Effective Address | 32/64       | Effective address of the SECS destination page.  |
| TMP_SECINFO    | Effective Address | 32/64       | Effective address of an SECINFO structure which contains security attributes of the SECS page to be added. |
| TMP_XSIZE      | SSA Size          | 64          | The size calculation of SSA frame.   |
| TMP_MISC_SIZE  | MISC Field Size   | 64          | Size of the selected MISC field components.  |
| TMPUPDATEFIELD | SHA256 Buffer     | 512         | Buffer used to hold data being added to TMP_SECS.MRENCLAVE.  |

IF (DS:RBX is not 32Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

TMP\_SRCPGE ← DS:RBX.SRCPGE;  
TMP\_SECINFO ← DS:RBX.SECINFO;

IF (DS:TMP\_SRCPGE is not 4KByte aligned or DS:TMP\_SECINFO is not 64Byte aligned)  
THEN #GP(0); FI;

IF (DS:RBX.LINADDR != 0 or DS:RBX.SECS ≠ 0)  
THEN #GP(0); FI;

(\* Check for misconfigured SECINFO flags\*)  
IF (DS:TMP\_SECINFO reserved fields are not zero or DS:TMP\_SECINFO.FLAGS.PT ≠ PT\_SECS) )  
THEN #GP(0); FI;

TMP\_SECS ← RCX;

IF (EPC entry in use)  
THEN  
IF (<<VMX non-root operation>> AND <<ENABLE\_EPC\_VIRTUALIZATION\_EXTENSIONS>>)  
THEN  
VMCS.Exit\_reason ← SGX\_CONFLICT;

## SGX INSTRUCTION REFERENCES

```

        VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_EXCEPTION;
        VMCS.Exit_qualification.error ← 0;
        VMCS.Guest-physical_address ←
            << translation of DS:TMP_SECS produced by paging >>;
        VMCS.Guest-linear_address ← DS:TMP_SECS;
    Deliver VMEXIT;
    ELSE
        #GP(0);
FI;

FI;

IF (EPC entry in use)
    THEN #GP(0); FI;

IF (EPCM(DS:RCX).VALID = 1)
    THEN #PF(DS:RCX); FI;

(* Copy 4KBytes from source page to EPC page*)
DS:RCX[32767:0] ← DS:TMP_SRCPAGE[32767:0];

(* Check lower 2 bits of XFRM are set *)
IF ( ( DS:TMP_SECS.ATTRIBUTES.XFRM BitwiseAND 03H) ≠ 03H)
    THEN #GP(0); FI;

IF (XFRM is illegal)
    THEN #GP(0); FI;

(* Make sure that the SECS does not have any unsupported MISCSELECT options*)
IF ( !(CPUID.(EAX=12H, ECX=0):EBX[31:0] & DS:TMP_SECS.MISCSELECT[31:0]) )
    THEN
        EPCM(DS:TMP_SECS).EntryLock.Release();
        #GP(0);
FI;

(* Compute size of MISC area *)
TMP_MISC_SIZE ← compute_misc_region_size();

(* Compute the size required to save state of the enclave on async exit, see Section 41.7.2.2*)
TMP_XSIZE ← compute_xsave_size(DS:TMP_SECS.ATTRIBUTES.XFRM) + GPR_SIZE + TMP_MISC_SIZE;

(* Ensure that the declared area is large enough to hold XSAVE and GPR stat *)
IF ( ( DS:TMP_SECS.SSAFRAMESIZE*4096 < TMP_XSIZE)
    THEN #GP(0); FI;

IF ( ( DS:TMP_SECS.ATTRIBUTES.MODE64BIT = 1) and (DS:TMP_SECS.BASEADDR is not canonical) )
    THEN #GP(0); FI;

IF ( ( DS:TMP_SECS.ATTRIBUTES.MODE64BIT = 0) and (DS:TMP_SECS.BASEADDR and 0FFFFFFF00000000H) )
    THEN #GP(0); FI;

IF ( ( DS:TMP_SECS.ATTRIBUTES.MODE64BIT = 0) and (DS:TMP_SECS.SIZE ≥ 2 ^ (CPUID.(EAX=12H, ECX=0):.EDX[7:0]) ) )
    THEN #GP(0); FI;

IF ( ( DS:TMP_SECS.ATTRIBUTES.MODE64BIT = 1) and (DS:TMP_SECS.SIZE ≥ 2 ^ (CPUID.(EAX=12H, ECX=0):.EDX[15:8]) ) )

```



```
THEN #GP(0); FI;
```

```
(* Enclave size must be at least 8192 bytes and must be power of 2 in bytes*)
```

```
IF (DS:TMP_SECS.SIZE < 8192 or popcnt(DS:TMP_SECS.SIZE) > 1)
```

```
THEN #GP(0); FI;
```

```
(* Ensure base address of an enclave is aligned on size*)
```

```
IF ( ( DS:TMP_SECS.BASEADDR and (DS:TMP_SECS.SIZE-1) )
```

```
THEN #GP(0); FI;
```

```
* Ensure the SECS does not have any unsupported attributes*)
```

```
IF ( ( DS:TMP_SECS.ATTRIBUTES and (~CR_SGX_ATTRIBUTES_MASK) )
```

```
THEN #GP(0); FI;
```

```
IF ( ( DS:TMP_SECS reserved fields are not zero)
```

```
THEN #GP(0); FI;
```

```
(* Verify that CONFIGID/CONFIGSVN are not set with attribute *)
```

```
IF ( ((DS:TMP_SECS.CONFIGID ≠ 0) or (DS:TMP_SECS.CONFIGSVN ≠ 0)) AND (DS:TMP_SECS.ATTRIBUTES.KSS == 0) )
```

```
THEN #GP(0); FI;
```

```
Clear DS:TMP_SECS to Uninitialized;
```

```
DS:TMP_SECS.MRENCLAVE ← SHA256INITIALIZE(DS:TMP_SECS.MRENCLAVE);
```

```
DS:TMP_SECS.ISVSVN ← 0;
```

```
DS:TMP_SECS.ISVPRODID ← 0;
```

```
(* Initialize hash updates etc*)
```

```
Initialize enclave's MRENCLAVE update counter;
```

```
(* Add "ECREATE" string and SECS fields to MRENCLAVE *)
```

```
TMPUPDATEFIELD[63:0] ← 0045544145524345H; // "ECREATE"
```

```
TMPUPDATEFIELD[95:64] ← DS:TMP_SECS.SSAFRAMESIZE;
```

```
TMPUPDATEFIELD[159:96] ← DS:TMP_SECS.SIZE;
```

```
TMPUPDATEFIELD[511:160] ← 0;
```

```
DS:TMP_SECS.MRENCLAVE ← SHA256UPDATE(DS:TMP_SECS.MRENCLAVE, TMPUPDATEFIELD)
```

```
INC enclave's MRENCLAVE update counter;
```

```
(* Set EID *)
```

```
DS:TMP_SECS.EID ← LockedXAdd(CR_NEXT_EID, 1);
```

```
(* Initialize the virtual child count to zero *)
```

```
DS:TMP_SECS.VIRTCHILDCNT ← 0;
```

```
(* Load ENCLAVECONTEXT with Address out of paging of SECS *)
```

```
<< store translation of DS:RCX produced by paging in SECS(DS:RCX).ENCLAVECONTEXT >>
```

```
(* Set the EPCM entry, first create SECS identifier and store the identifier in EPCM *)
```

```
EPCM(DS:TMP_SECS).PT ← PT_SECS;
```

```
EPCM(DS:TMP_SECS).ENCLAVEADDRESS ← 0;
```

```
EPCM(DS:TMP_SECS).R ← 0;
```

```
EPCM(DS:TMP_SECS).W ← 0;
```

```
EPCM(DS:TMP_SECS).X ← 0;
```

```
(* Set EPCM entry fields *)
```

## SGX INSTRUCTION REFERENCES

EPCM(DS:RCX).BLOCKED ← 0;  
EPCM(DS:RCX).PENDING ← 0;  
EPCM(DS:RCX).MODIFIED ← 0;  
EPCM(DS:RCX).PR ← 0;  
EPCM(DS:RCX).VALID ← 1;

### Flags Affected

None

### Protected Mode Exceptions

#GP(0)            If a memory operand effective address is outside the DS segment limit.  
                  If a memory operand is not properly aligned.  
                  If the reserved fields are not zero.  
                  If PAGEINFO.SECS is not zero.  
                  If PAGEINFO.LINADDR is not zero.  
                  If the SECS destination is locked.  
                  If SECS.SSAFRAMESIZE is insufficient.

#PF(error code)    If a page fault occurs in accessing memory operands.  
                  If the SECS destination is outside the EPC.

### 64-Bit Mode Exceptions

#GP(0)            If a memory address is non-canonical form.  
                  If a memory operand is not properly aligned.  
                  If the reserved fields are not zero.  
                  If PAGEINFO.SECS is not zero.  
                  If PAGEINFO.LINADDR is not zero.  
                  If the SECS destination is locked.  
                  If SECS.SSAFRAMESIZE is insufficient.

#PF(error code)    If a page fault occurs in accessing memory operands.  
                  If the SECS destination is outside the EPC.

## EDBGRD—Read From a Debug Enclave

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 04H<br>ENCLS[EDBGRD] | IR    | V/V                          | SGX1                     | This leaf function reads a dword/quadword from a debug enclave. |

### Instruction Operand Encoding

| Op/En | EAX         | RBX                                  | RCX                                      |
|-------|-------------|--------------------------------------|--|
| IR    | EDBGRD (In) | Data read from a debug enclave (Out) | Address of source memory in the EPC (In) |

### Description

This leaf function copies a quadword/doubleword from an EPC page belonging to a debug enclave into the RBX register. Eight bytes are read in 64-bit mode, four bytes are read in non-64-bit modes. The size of data read cannot be overridden.

The effective address of the source location inside the EPC is provided in the register RCX.

### EDBGRD Memory Parameter Semantics

|                                  |
|----------------------------------|
| EPCQW                            |
| Read access permitted by Enclave |

The error codes are:

**Table 40-17. EDBGRD Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EDBGRD successful.  |
| SGX_PAGE_NOT_DEBUGGABLE     | The EPC page cannot be accessed because it is in the PENDING or MODIFIED state. |

The instruction faults if any of the following:

### EDBGRD Faulting Conditions

|  |   |
|--|---|
| RCX points into a page that is an SECS.  | RCX does not resolve to a naturally aligned linear address.   |
| RCX points to a page that does not belong to an enclave that is in debug mode. | RCX points to a location inside a TCS that is beyond the architectural size of the TCS (SGX_TCS_LIMIT). |
| An operand causing any segment violation.                                      | May page fault.   |
| CPL > 0.   |   |

This instruction ignores the EPCM RWX attributes on the enclave page. Consequently, violation of EPCM RWX attributes via EDBGRD does not result in a #GP.

Concurrency Restrictions

**Table 40-18. Base Concurrency Restrictions of EDBGRD**

| Leaf   | Parameter       | Base Concurrency Restrictions |             |                                    |
|--------|-----------------|-------------------------------|-------------|------------------------------------|
|        |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EDBGRD | Target [DS:RCX] | Shared                        | #GP         |                                    |

**Table 40-19. Additional Concurrency Restrictions of EDBGRD**

| Leaf   | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|        |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EDBGRD | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

**Temp Variables in EDBGRD Operational Flow**

| Name       | Type   | Size (Bits) | Description  |
|------------|--------|-------------|--|
| TMP_MODE64 | Binary | 1           | ((IA32_EFER.LMA = 1) && (CS.L = 1))                                      |
| TMP_SECS   |        | 64          | Physical address of SECS of the enclave to which source operand belongs. |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

IF ( (TMP\_MODE64 = 1) and (DS:RCX is not 8Byte Aligned) )  
 THEN #GP(0); FI;

IF ( (TMP\_MODE64 = 0) and (DS:RCX is not 4Byte Aligned) )  
 THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

(\* make sure no other Intel SGX instruction is accessing EPCM \*)

IF (Other EPCM modifying instructions executing)  
 THEN #GP(0); FI;

IF (EPCM(DS:RCX). VALID = 0)  
 THEN #PF(DS:RCX); FI;

(\* make sure that DS:RCX (SOURCE) is pointing to a PT\_REG or PT\_TCS or PT\_VA \*)

IF ( (EPCM(DS:RCX).PT ≠ PT\_REG) and (EPCM(DS:RCX).PT ≠ PT\_TCS) and (EPCM(DS:RCX).PT ≠ PT\_VA) )  
 THEN #PF(DS:RCX); FI;

(\* make sure that DS:RCX points to an accessible EPC page \*)

IF ( (EPCM(DS:RCX).PENDING is not 0 or (EPCM(DS:RCX).MODIFIED is not 0) )  
 THEN  
 RFLAGS.ZF ← 1;  
 RAX ← SGX\_PAGE\_NOT\_DEBUGGABLE;

```

    GOTO DONE;
FI;

(* If source is a TCS, then make sure that the offset into the page is not beyond the TCS size*)
IF ( ( EPCM(DS:RCX).PT = PT_TCS) and ((DS:RCX) & FFFH ≥ SGX_TCS_LIMIT) )
    THEN #GP(0); FI;

(* make sure the enclave owning the PT_REG or PT_TCS page allow debug *)
IF ( (EPCM(DS:RCX).PT = PT_REG) or (EPCM(DS:RCX).PT = PT_TCS) )
    THEN
        TMP_SECS ← GET_SECS_ADDRESS;
        IF (TMP_SECS.ATTRIBUTES.DEBUG = 0)
            THEN #GP(0); FI;
        IF ( (TMP_MODE64 = 1) )
            THEN RBX[63:0] ← (DS:RCX)[63:0];
            ELSE EBX[31:0] ← (DS:RCX)[31:0];
        FI;
    ELSE
        TMP_64BIT_VAL[63:0] ← (DS:RCX)[63:0] & (~07H); // Read contents from VA slot
        IF (TMP_MODE64 = 1)
            THEN
                IF (TMP_64BIT_VAL ≠ 0H)
                    THEN RBX[63:0] ← 0FFFFFFFFFFFFFFFFH;
                    ELSE RBX[63:0] ← 0H;
                FI;
            ELSE
                IF (TMP_64BIT_VAL ≠ 0H)
                    THEN EBX[31:0] ← 0FFFFFFFFH;
                    ELSE EBX[31:0] ← 0H;
                FI;
        FI;

(* clear EAX and ZF to indicate successful completion *)
RAX ← 0;
RFLAGS.ZF ← 0;

DONE:
(* clear flags *)
RFLAGS.CF,PF,AF,OF,SF ← 0;

```

### Flags Affected

None

### Protected Mode Exceptions

|        |  |
|--------|--|
| #GP(0) | <p>If the address in RCS violates DS limit or access rights.</p> <p>If DS segment is unusable.</p> <p>If RCX points to a memory location not 4Byte-aligned.</p> <p>If the address in RCX points to a page belonging to a non-debug enclave.</p> <p>If the address in RCX points to a page which is not PT_TCS, PT_REG or PT_VA.</p> <p>If the address in RCX points to a location inside TCS that is beyond SGX_TCS_LIMIT.</p> |
|--------|--|

## SGX INSTRUCTION REFERENCES

#PF(error code)    If a page fault occurs in accessing memory operands.  
                         If the address in RCX points to a non-EPC page.  
                         If the address in RCX points to an invalid EPC page.

### 64-Bit Mode Exceptions

#GP(0)              If RCX is non-canonical form.  
                         If RCX points to a memory location not 8Byte-aligned.  
                         If the address in RCX points to a page belonging to a non-debug enclave.  
                         If the address in RCX points to a page which is not PT\_TCS, PT\_REG or PT\_VA.  
                         If the address in RCX points to a location inside TCS that is beyond SGX\_TCS\_LIMIT.

#PF(error code)    If a page fault occurs in accessing memory operands.  
                         If the address in RCX points to a non-EPC page.  
                         If the address in RCX points to an invalid EPC page.

## EDBGWR—Write to a Debug Enclave

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|----------------------------|-------|------------------------------|--------------------------|--|
| EAX = 05H<br>ENCLS[EDBGWR] | IR    | V/V                          | SGX1                     | This leaf function writes a dword/quadword to a debug enclave. |

### Instruction Operand Encoding

| Op/En | EAX         | RBX  | RCX                                      |
|-------|-------------|--|--|
| IR    | EDBGWR (In) | Data to be written to a debug enclave (In) | Address of Target memory in the EPC (In) |

### Description

This leaf function copies the content in EBX/RBX to an EPC page belonging to a debug enclave. Eight bytes are written in 64-bit mode, four bytes are written in non-64-bit modes. The size of data cannot be overridden.

The effective address of the source location inside the EPC is provided in the register RCX.

### EDBGWR Memory Parameter Semantics

| EPCQW                             |
|-----------------------------------|
| Write access permitted by Enclave |

The instruction faults if any of the following:

### EDBGWR Faulting Conditions

|  |   |
|--|---|
| RCX points into a page that is an SECS.  | RCX does not resolve to a naturally aligned linear address.       |
| RCX points to a page that does not belong to an enclave that is in debug mode. | RCX points to a location inside a TCS that is not the FLAGS word. |
| An operand causing any segment violation.                                      | May page fault.   |
| CPL > 0.   |   |

The error codes are:

**Table 40-20. EDBGWR Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EDBGWR successful.  |
| SGX_PAGE_NOT_DEBUGGABLE     | The EPC page cannot be accessed because it is in the PENDING or MODIFIED state. |

This instruction ignores the EPCM RWX attributes on the enclave page. Consequently, violation of EPCM RWX attributes via EDBGWR does not result in a #GP.

### Concurrency Restrictions

**Table 40-21. Base Concurrency Restrictions of EDBGWR**

| Leaf   | Parameter       | Base Concurrency Restrictions |             |                                    |
|--------|-----------------|-------------------------------|-------------|------------------------------------|
|        |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EDBGWR | Target [DS:RCX] | Shared                        | #GP         |                                    |

**Table 40-22. Additional Concurrency Restrictions of EDBGWR**

| Leaf   | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|        |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EDBGWR | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EDBGWR Operational Flow**

| Name       | Type   | Size (Bits) | Description  |
|------------|--------|-------------|--|
| TMP_MODE64 | Binary | 1           | ((IA32_EFER.LMA = 1) && (CS.L = 1)).                                     |
| TMP_SECS   |        | 64          | Physical address of SECS of the enclave to which source operand belongs. |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

IF ( (TMP\_MODE64 = 1) and (DS:RCX is not 8Byte Aligned) )  
 THEN #GP(0); FI;

IF ( (TMP\_MODE64 = 0) and (DS:RCX is not 4Byte Aligned) )  
 THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

(\* make sure no other Intel SGX instruction is accessing EPCM \*)  
 IF (Other EPCM modifying instructions executing)  
 THEN #GP(0); FI;

IF (EPCM(DS:RCX).VALID = 0)  
 THEN #PF(DS:RCX); FI;

(\* make sure that DS:RCX (DST) is pointing to a PT\_REG or PT\_TCS \*)  
 IF ( (EPCM(DS:RCX).PT ≠ PT\_REG) and (EPCM(DS:RCX).PT ≠ PT\_TCS) )  
 THEN #PF(DS:RCX); FI;

(\* make sure that DS:RCX points to an accessible EPC page \*)  
 IF ( (EPCM(DS:RCX).PENDING is not 0) or (EPCM(DS:RCX).MODIFIED is not 0) )  
 THEN  
     RFLAGS.ZF ← 1;  
     RAX ← SGX\_PAGE\_NOT\_DEBUGGABLE;  
     GOTO DONE;

FI;

(\* If destination is a TCS, then make sure that the offset into the page can only point to the FLAGS field\*)  
 IF ( (EPCM(DS:RCX).PT = PT\_TCS) and ((DS:RCX) & FF8H ≠ offset\_of\_FLAGS & 0FF8H) )  
 THEN #GP(0); FI;

(\* Locate the SECS for the enclave to which the DS:RCX page belongs \*)



```
TMP_SECS ← GET_SECS_PHYS_ADDRESS(EPCM(DS:RCX).ENCLAVESECS);
```

```
(* make sure the enclave owning the PT_REG or PT_TCS page allow debug *)
```

```
IF (TMP_SECS.ATTRIBUTES.DEBUG = 0)
```

```
    THEN #GP(0); FI;
```

```
IF ( (TMP_MODE64 = 1) )
```

```
    THEN (DS:RCX)[63:0] ← RBX[63:0];
```

```
    ELSE (DS:RCX)[31:0] ← EBX[31:0];
```

```
FI;
```

```
(* clear EAX and ZF to indicate successful completion *)
```

```
RAX ← 0;
```

```
RFLAGS.ZF ← 0;
```

```
DONE:
```

```
(* clear flags *)
```

```
RFLAGS.CF,PF,AF,OF,SF ← 0
```

### Flags Affected

None

### Protected Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <p>If the address in RCS violates DS limit or access rights.</p> <p>If DS segment is unusable.</p> <p>If RCX points to a memory location not 4Byte-aligned.</p> <p>If the address in RCX points to a page belonging to a non-debug enclave.</p> <p>If the address in RCX points to a page which is not PT_TCS or PT_REG.</p> <p>If the address in RCX points to a location inside TCS that is not the FLAGS word.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory operands.</p> <p>If the address in RCX points to a non-EPC page.</p> <p>If the address in RCX points to an invalid EPC page.</p>  |

### 64-Bit Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <p>If RCX is non-canonical form.</p> <p>If RCX points to a memory location not 8Byte-aligned.</p> <p>If the address in RCX points to a page belonging to a non-debug enclave.</p> <p>If the address in RCX points to a page which is not PT_TCS or PT_REG.</p> <p>If the address in RCX points to a location inside TCS that is not the FLAGS word.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory operands.</p> <p>If the address in RCX points to a non-EPC page.</p> <p>If the address in RCX points to an invalid EPC page.</p>  |

## EEXTEND—Extend Uninitialized Enclave Measurement by 256 Bytes

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 06H<br>ENCLS[EEXTEND] | IR    | V/V                          | SGX1                     | This leaf function measures 256 bytes of an uninitialized enclave page. |

### Instruction Operand Encoding

| Op/En | EAX          | EBX  | RCX   |
|-------|--------------|--|---|
| IR    | EEXTEND (In) | Effective address of the SECS of the data chunk (In) | Effective address of a 256-byte chunk in the EPC (In) |

### Description

This leaf function updates the MRENCLAVE measurement register of an SECS with the measurement of an EXTEND string comprising of “EEXTEND” || ENCLAVEOFFSET || PADDING || 256 bytes of the enclave page. This instruction can only be executed when current privilege level is 0 and the enclave is uninitialized.

RBX contains the effective address of the SECS of the region to be measured. The address must be the same as the one used to add the page into the enclave.

RCX contains the effective address of the 256 byte region of an EPC page to be measured. The DS segment is used to create linear addresses. Segment override is not supported.

### EEXTEND Memory Parameter Semantics

|                        |
|------------------------|
| EPC[RCX]               |
| Read access by Enclave |

The instruction faults if any of the following:

### EEXTEND Faulting Conditions

|   |  |
|---|--|
| RBX points to an address not 4KBytes aligned. | RBX does not resolve to an SECS.                       |
| RBX does not point to an SECS page.           | RBX does not point to the SECS page of the data chunk. |
| RCX points to an address not 256B aligned.    | RCX points to an unused page or a SECS.                |
| RCX does not resolve in an EPC page.          | If SECS is locked.                                     |
| If the SECS is already initialized.           | May page fault.  |
| CPL > 0.                                      |  |

### Concurrency Restrictions

**Table 40-23. Base Concurrency Restrictions of EEXTEND**

| Leaf    | Parameter       | Base Concurrency Restrictions |             |                                    |
|---------|-----------------|-------------------------------|-------------|------------------------------------|
|         |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EEXTEND | Target [DS:RCX] | Shared                        | #GP         |                                    |
|         | SECS [DS:RBX]   | Concurrent                    |             |                                    |

Table 40-24. Additional Concurrency Restrictions of EEXTEND

| Leaf    | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EEXTEND | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|         | SECS [DS:RBX]   | Concurrent                                      |             | Exclusive                | #GP         | Concurrent          |             |

## Operation

### Temp Variables in EEXTEND Operational Flow

| Name                  | Type           | Size (Bits) | Description  |
|-----------------------|----------------|-------------|--|
| TMP_SECS              |                | 64          | Physical address of SECS of the enclave to which source operand belongs. |
| TMP_ENCLAVEOFFS<br>ET | Enclave Offset | 64          | The page displacement from the enclave base address.                     |
| TMPUPDATEFIELD        | SHA256 Buffer  | 512         | Buffer used to hold data being added to TMP_SECS.MRENCLAVE.              |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

IF (DS:RBX is not 4096 Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RBX does resolve to an EPC page)  
THEN #PF(DS:RBX); FI;

IF (DS:RCX is not 256Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

(\* make sure no other Intel SGX instruction is accessing EPCM \*)  
IF (Other instructions accessing EPCM)  
THEN #GP(0); FI;

IF (EPCM(DS:RCX).VALID = 0)  
THEN #PF(DS:RCX); FI;

(\* make sure that DS:RCX (DST) is pointing to a PT\_REG or PT\_TCS \*)  
IF ( (EPCM(DS:RCX).PT ≠ PT\_REG) and (EPCM(DS:RCX).PT ≠ PT\_TCS) )  
THEN #PF(DS:RCX); FI;

TMP\_SECS ← Get\_SECS\_ADDRESS();

IF (DS:RBX does not resolve to TMP\_SECS)  
THEN #GP(0); FI;

(\* make sure no other instruction is accessing MRENCLAVE or ATTRIBUETS.INIT \*)  
IF ( (Other instruction accessing MRENCLAVE) or (Other instructions checking or updating the initialized state of the SECS) )  
THEN #GP(0); FI;

(\* Calculate enclave offset \*)

TMP\_ENCLAVEOFFSET ← EPCM(DS:RCX).ENCLAVEADDRESS - TMP\_SECS.BASEADDR;  
 TMP\_ENCLAVEOFFSET ← TMP\_ENCLAVEOFFSET + (DS:RCX & 0FFFH)

(\* Add EEXTEND message and offset to MRENCLAVE \*)

TMPUPDATEFIELD[63:0] ← 00444E4554584545H; // "EEXTEND"  
 TMPUPDATEFIELD[127:64] ← TMP\_ENCLAVEOFFSET;  
 TMPUPDATEFIELD[511:128] ← 0; // 48 bytes  
 TMP\_SECS.MRENCLAVE ← SHA256UPDATE(TMP\_SECS.MRENCLAVE, TMPUPDATEFIELD)  
 INC enclave's MRENCLAVE update counter;

(\*Add 256 bytes to MRENCLAVE, 64 byte at a time \*)

TMP\_SECS.MRENCLAVE ← SHA256UPDATE(TMP\_SECS.MRENCLAVE, DS:RCX[511:0]);  
 TMP\_SECS.MRENCLAVE ← SHA256UPDATE(TMP\_SECS.MRENCLAVE, DS:RCX[1023: 512]);  
 TMP\_SECS.MRENCLAVE ← SHA256UPDATE(TMP\_SECS.MRENCLAVE, DS:RCX[1535: 1024]);  
 TMP\_SECS.MRENCLAVE ← SHA256UPDATE(TMP\_SECS.MRENCLAVE, DS:RCX[2047: 1536]);  
 INC enclave's MRENCLAVE update counter by 4;

### Flags Affected

None

### Protected Mode Exceptions

#GP(0) If the address in RBX is outside the DS segment limit.  
 If RBX points to an SECS page which is not the SECS of the data chunk.  
 If the address in RCX is outside the DS segment limit.  
 If RCX points to a memory location not 256Byte-aligned.  
 If another instruction is accessing MRENCLAVE.  
 If another instruction is checking or updating the SECS.  
 If the enclave is already initialized.

#PF(error code) If a page fault occurs in accessing memory operands.  
 If the address in RBX points to a non-EPC page.  
 If the address in RCX points to a page which is not PT\_TCS or PT\_REG.  
 If the address in RCX points to a non-EPC page.  
 If the address in RCX points to an invalid EPC page.

### 64-Bit Mode Exceptions

#GP(0) If RBX is non-canonical form.  
 If RBX points to an SECS page which is not the SECS of the data chunk.  
 If RCX is non-canonical form.  
 If RCX points to a memory location not 256 Byte-aligned.  
 If another instruction is accessing MRENCLAVE.  
 If another instruction is checking or updating the SECS.  
 If the enclave is already initialized.

#PF(error code) If a page fault occurs in accessing memory operands.  
 If the address in RBX points to a non-EPC page.  
 If the address in RCX points to a page which is not PT\_TCS or PT\_REG.  
 If the address in RCX points to a non-EPC page.  
 If the address in RCX points to an invalid EPC page.

## EINIT—Initialize an Enclave for Execution

| Opcode/<br>Instruction    | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|---------------------------|-------|------------------------------|--------------------------|--|
| EAX = 02H<br>ENCLS[EINIT] | IR    | V/V                          | SGX1                     | This leaf function initializes the enclave and makes it ready to execute enclave code. |

### Instruction Operand Encoding

| Op/En | EAX        |                  | RBX                       | RCX                  | RDX                        |
|-------|------------|------------------|---------------------------|----------------------|----------------------------|
| IR    | EINIT (In) | Error code (Out) | Address of SIGSTRUCT (In) | Address of SECS (In) | Address of EINITTOKEN (In) |

### Description

This leaf function is the final instruction executed in the enclave build process. After EINIT, the MRENCLAVE measurement is complete, and the enclave is ready to start user code execution using the EENTER instruction.

EINIT takes the effective address of a SIGSTRUCT and EINITTOKEN. The SIGSTRUCT describes the enclave including MRENCLAVE, ATTRIBUTES, ISVSVN, a 3072 bit RSA key, and a signature using the included key. SIGSTRUCT must be populated with two values, q1 and q2. These are calculated using the formulas shown below:

$$q1 = \text{floor}(\text{Signature}^2 / \text{Modulus});$$

$$q2 = \text{floor}((\text{Signature}^3 - q1 * \text{Signature} * \text{Modulus}) / \text{Modulus});$$

The EINITTOKEN contains the MRENCLAVE, MRSIGNER, and ATTRIBUTES. These values must match the corresponding values in the SECS. If the EINITTOKEN was created with a debug launch key, the enclave must be in debug mode as well.




Figure 40-1. Relationships Between SECS, SIGSTRUCT and EINITTOKEN

**EINIT Memory Parameter Semantics**

|                       |                              |                       |
|-----------------------|------------------------------|-----------------------|
| SIGSTRUCT             | SECS                         | EINITOKEN             |
| Access by non-Enclave | Read/Write access by Enclave | Access by non-Enclave |

EINIT performs the following steps, which can be seen in Figure 40-1:

Validates that SIGSTRUCT is signed using the enclosed public key.

Checks that the completed computation of SECS.MRENCLAVE equals SIGSTRUCT.HASHENCLAVE.

Checks that no reserved bits are set to 1 in SIGSTRUCT.ATTRIBUTES and no reserved bits in SIGSTRUCT.ATTRIBUTESMASK are set to 0.

Checks that no controlled ATTRIBUTES bits are set in SIGSTRUCT.ATTRIBUTES unless the SHA256 digest of SIGSTRUCT.MODULUS equals IA32\_SGX\_LEPUBKEYHASH.

Checks that SIGSTRUCT.ATTRIBUTES equals the result of logically and-ing SIGSTRUCT.ATTRIBUTESMASK with SECS.ATTRIBUTES.

If EINITOKEN.VALID is 0, checks that the SHA256 digest of SIGSTRUCT.MODULUS equals IA32\_SGX\_LEPUBKEYHASH.

If EINITOKEN.VALID is 1, checks the validity of EINITOKEN.

If EINITOKEN.VALID is 1, checks that EINITOKEN.MRENCLAVE equals SECS.MRENCLAVE.

If EINITOKEN.VALID is 1 and EINITOKEN.ATTRIBUTES.DEBUG is 1, SECS.ATTRIBUTES.DEBUG must be 1.

Commits SECS.MRENCLAVE, and sets SECS.MRSIGNER, SECS.ISVSVN, and SECS.ISVPRODID based on SIGSTRUCT.

Update the SECS as Initialized.

Periodically, EINIT polls for certain asynchronous events. If such an event is detected, it completes with failure code (ZF=1 and RAX = SGX\_UNMASKED\_EVENT), and RIP is incremented to point to the next instruction. These events includes external interrupts, non-maskable interrupts, system-management interrupts, machine checks, INIT signals, and the VMX-preemption timer. EINIT does not fail if the pending event is inhibited (e.g., external interrupts could be inhibited due to blocking by MOV SS blocking or by STI).

The following bits in RFLAGS are cleared: CF, PF, AF, OF, and SF. When the instruction completes with an error, RFLAGS.ZF is set to 1, and the corresponding error bit is set in RAX. If no error occurs, RFLAGS.ZF is cleared and RAX is set to 0.

The error codes are:

**Table 40-25. EINIT Return Value in RAX**

| Error Code (see Table 40-4) | Description  |
|-----------------------------|--|
| No Error                    | EINIT successful.  |
| SGX_INVALID_SIG_STRUCT      | If SIGSTRUCT contained an invalid value.   |
| SGX_INVALID_ATTRIBUTE       | If SIGSTRUCT contains an unauthorized attributes mask.   |
| SGX_INVALID_MEASUREMENT     | If SIGSTRUCT contains an incorrect measurement.<br>If EINITOKEN contains an incorrect measurement. |
| SGX_INVALID_SIGNATURE       | If signature does not validate with enclosed public key.   |
| SGX_INVALID_LICENSE         | If license is invalid.   |
| SGX_INVALID_CPUSVN          | If license SVN is unsupported.   |
| SGX_UNMASKED_EVENT          | If an unmasked event is received before the instruction completes its operation.                   |

## Concurrency Restrictions

Table 40-26. Base Concurrency Restrictions of EINIT

| Leaf  | Parameter     | Base Concurrency Restrictions |             |                                    |
|-------|---------------|-------------------------------|-------------|------------------------------------|
|       |               | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EINIT | SECS [DS:RCX] | Shared                        | #GP         |                                    |

Table 40-27. Additional Concurrency Restrictions of ENIT

| Leaf  | Parameter     | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------|---------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|       |               | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|       |               | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EINIT | SECS [DS:RCX] | Concurrent                                      |             | Exclusive                | #GP         | Concurrent          |             |

## Operation

## Temp Variables in EINIT Operational Flow

| Name                  | Type                | Size      | Description  |
|-----------------------|---------------------|-----------|--|
| TMP_SIG               | SIGSTRUCT           | 1808Bytes | Temp space for SIGSTRUCT.  |
| TMP_TOKEN             | EINITTOKEN          | 304Bytes  | Temp space for EINITTOKEN.   |
| TMP_MRENCLAVE         |                     | 32Bytes   | Temp space for calculating MRENCLAVE.  |
| TMP_MRSIGNER          |                     | 32Bytes   | Temp space for calculating MRSIGNER.   |
| CONTROLLED_ATTRIBUTES | ATTRIBUTES          | 16Bytes   | Constant mask of all ATTRIBUTE bits that can only be set for authorized enclaves.              |
| TMP_KEYDEPENDENCIES   | Buffer              | 224Bytes  | Temp space for key derivation.   |
| TMP_EINITTOKENKEY     |                     | 16Bytes   | Temp space for the derived EINITTOKEN Key.   |
| TMP_SIG_PADDING       | PKCS Padding Buffer | 352Bytes  | The value of the top 352 bytes from the computation of Signature <sup>3</sup> modulo MRSIGNER. |

(\* make sure SIGSTRUCT and SECS are aligned \*)

IF ( (DS:RBX is not 4KByte Aligned) or (DS:RCX is not 4KByte Aligned) )  
THEN #GP(0); FI;

(\* make sure the EINITTOKEN is aligned \*)

IF (DS:RDX is not 512Byte Aligned)  
THEN #GP(0); FI;

(\* make sure the SECS is inside the EPC \*)

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

TMP\_SIG[14463:0] ← DS:RBX[14463:0]; // 1808 bytes

TMP\_TOKEN[2423:0] ← DS:RDX[2423:0]; // 304 bytes

## SGX INSTRUCTION REFERENCES

(\* Verify SIGSTRUCT Header. \*)

```
IF ( (TMP_SIG.HEADER ≠ 06000000E10000000000010000000000h) or
    ((TMP_SIG.VENDOR ≠ 0) and (TMP_SIG.VENDOR ≠ 00008086h) ) or
    (TMP_SIG.HEADER2 ≠ 01010000600000006000000001000000h) or
    (TMP_SIG.EXPONENT ≠ 00000003h) or (Reserved space is not 0's) )
    THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_SIG_STRUCT;
    GOTO EXIT;
FI;
```

(\* Open "Event Window" Check for Interrupts. Verify signature using embedded public key, q1, and q2. Save upper 352 bytes of the PKCS1.5 encoded message into the TMP\_SIG\_PADDING\*)

```
IF (interrupt was pending) THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_UNMASKED_EVENT;
    GOTO EXIT;
FI
```

```
IF (signature failed to verify) THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_SIGNATURE;
    GOTO EXIT;
FI;
```

(\*Close "Event Window" \*)

(\* make sure no other Intel SGX instruction is modifying SECS\*)

```
IF (Other instructions modifying SECS)
    THEN #GP(0); FI;
```

```
IF ( (EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PT ≠ PT_SECS) )
    THEN #PF(DS:RCX); FI;
```

(\* Verify ISVFAMILYID is not used on an enclave with KSS disabled \*)

```
IF ((TMP_SIG.ISVFAMILYID != 0) AND (DS:RCX.ATTRIBUTES.KSS == 0))
    THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_SIG_STRUCT;
    GOTO EXIT;
FI;
```

(\* make sure no other instruction is accessing MRENCLAVE or ATTRIBUTES.INIT \*)

```
IF ( (Other instruction modifying MRENCLAVE) or (Other instructions modifying the SECS's Initialized state))
    THEN #GP(0); FI;
```

(\* Calculate finalized version of MRENCLAVE \*)

(\* SHA256 algorithm requires one last update that compresses the length of the hashed message into the output SHA256 digest \*)

```
TMP_ENCLAVE ← SHA256FINAL( (DS:RCX).MRENCLAVE, enclave's MRENCLAVE update count *512);
```

(\* Verify MRENCLAVE from SIGSTRUCT \*)

```
IF (TMP_SIG.ENCLAVEHASH ≠ TMP_MRENCLAVE)
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_MEASUREMENT;
    GOTO EXIT;
FI;
```



```
TMP_MRSIGNER ← SHA256(TMP_SIG.MODULUS)
```

```
(* if controlled ATTRIBUTES are set, SIGSTRUCT must be signed using an authorized key *)
```

```
CONTROLLED_ATTRIBUTES ← 000000000000020H;
```

```
IF ( ( (DS:RCX.ATTRIBUTES & CONTROLLED_ATTRIBUTES) ≠ 0) and (TMP_MRSIGNER ≠ IA32_SGXLEPUBKEYHASH) )
```

```
    RFLAGS.ZF ← 1;
```

```
    RAX ← SGX_INVALID_ATTRIBUTE;
```

```
    GOTO EXIT;
```

```
FI;
```

```
(* Verify SIGSTRUCT.ATTRIBUTE requirements are met *)
```

```
IF ( (DS:RCX.ATTRIBUTES & TMP_SIG.ATTRIBUTEMASK) ≠ (TMP_SIG.ATTRIBUTE & TMP_SIG.ATTRIBUTEMASK) )
```

```
    RFLAGS.ZF ← 1;
```

```
    RAX ← SGX_INVALID_ATTRIBUTE;
```

```
    GOTO EXIT;
```

```
FI;
```

```
(*Verify SIGSTRUCT.MISCSELECT requirements are met *)
```

```
IF ( (DS:RCX.MISCSELECT & TMP_SIG.MISCMASK) ≠ (TMP_SIG.MISCSELECT & TMP_SIG.MISCMASK) )
```

```
    THEN
```

```
        RFLAGS.ZF ← 1;
```

```
        RAX ← SGX_INVALID_ATTRIBUTE;
```

```
    GOTO EXIT
```

```
FI;
```

```
(* if EINITOKEN.VALID[0] is 0, verify the enclave is signed by an authorized key *)
```

```
IF (TMP_TOKEN.VALID[0] = 0)
```

```
    IF (TMP_MRSIGNER ≠ IA32_SGXLEPUBKEYHASH)
```

```
        RFLAGS.ZF ← 1;
```

```
        RAX ← SGX_INVALID_EINITOKEN;
```

```
        GOTO EXIT;
```

```
    FI;
```

```
    GOTO COMMIT;
```

```
FI;
```

```
(* Debug Launch Enclave cannot launch Production Enclaves *)
```

```
IF ( (DS:RDX.MASKEDATTRIBUTESLE.DEBUG = 1) and (DS:RCX.ATTRIBUTES.DEBUG = 0) )
```

```
    RFLAGS.ZF ← 1;
```

```
    RAX ← SGX_INVALID_EINITOKEN;
```

```
    GOTO EXIT;
```

```
FI;
```

```
(* Check reserve space in EINIT token includes reserved regions and upper bits in valid field *)
```

```
IF (TMP_TOKEN reserved space is not clear)
```

```
    RFLAGS.ZF ← 1;
```

```
    RAX ← SGX_INVALID_EINITOKEN;
```

```
    GOTO EXIT;
```

```
FI;
```

```
(* EINIT token must be ≤ CR_CPUSVN *)
```

```
IF (TMP_TOKEN.CPUSVN > CR_CPUSVN)
```

```
    RFLAGS.ZF ← 1;
```

```
    RAX ← SGX_INVALID_CPUSVN;
```

## SGX INSTRUCTION REFERENCES

```
GOTO EXIT;
FI;
```

```
(* Derive Launch key used to calculate EINITTOKEN.MAC *)
```

```
HARDCODED_PKCS1_5_PADDING[15:0] ← 0100H;
HARDCODED_PKCS1_5_PADDING[2655:16] ← SignExtend330Byte(-1); // 330 bytes of 0FFH
HARDCODED_PKCS1_5_PADDING[2815:2656] ← 2004000501020403650148866009060D30313000H;
```

```
TMP_KEYDEPENDENCIES.KEYNAME ← EINITTOKEN_KEY;
TMP_KEYDEPENDENCIES.ISVFAMILYID ← 0;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVPRODID ← TMP_TOKEN.ISVPRODIDLE;
TMP_KEYDEPENDENCIES.ISVSVN ← TMP_TOKEN.ISVSVN;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← CR_SGXOWNEREPOCH;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_TOKEN.MASKEDATTRIBUTESLE;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← 0;
TMP_KEYDEPENDENCIES.MRENCLAVE ← 0;
TMP_KEYDEPENDENCIES.MRSIGNER ← IA32_SGXLEPUBKEYHASH;
TMP_KEYDEPENDENCIES.KEYID ← TMP_TOKEN.KEYID;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← TMP_TOKEN.CPUSVN;
TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_TOKEN.MASKEDMISCSELECTLE;
TMP_KEYDEPENDENCIES.MISCMASK ← 0;
TMP_KEYDEPENDENCIES.PADDING ← HARDCODED_PKCS1_5_PADDING;
TMP_KEYDEPENDENCIES.KEYPOLICY ← 0;
TMP_KEYDEPENDENCIES.CONFIGID ← 0;
TMP_KEYDEPENDENCIES.CONFIGSVN ← 0;
```

```
(* Calculate the derived key*)
```

```
TMP_EINITTOKENKEY ← derivekey(TMP_KEYDEPENDENCIES);
```

```
(* Verify EINITTOKEN was generated using this CPU's Launch key and that it has not been modified since issuing by the Launch Enclave. Only 192 bytes of EINITTOKEN are CMACed *)
```

```
IF (TMP_TOKEN.MAC ≠ CMAC(TMP_EINITTOKENKEY, TMP_TOKEN[1535:0] ))
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_EINITTOKEN;
    GOTO EXIT;
```

```
FI;
```

```
(* Verify EINITTOKEN (RDX) is for this enclave *)
```

```
IF (TMP_TOKEN.MRENCLAVE ≠ TMP_MRENCLAVE) or (TMP_TOKEN.MRSIGNER ≠ TMP_MRSIGNER) )
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_MEASUREMENT;
    GOTO EXIT;
```

```
FI;
```

```
(* Verify ATTRIBUTES in EINITTOKEN are the same as the enclave's *)
```

```
IF (TMP_TOKEN.ATTRIBUTES ≠ DS:RCX.ATTRIBUTES)
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_EINIT_ATTRIBUTE;
    GOTO EXIT;
```

```
FI;
```

```
COMMIT:
```

(\* Commit changes to the SECS; Set ISVPRODID, ISVSVN, MRSIGNER, INIT ATTRIBUTE fields in SECS (RCX) \*)

DS:RCX.MRENCLAVE ← TMP\_MRENCLAVE;

(\* MRSIGNER stores a SHA256 in little endian implemented natively on x86 \*)

DS:RCX.MRSIGNER ← TMP\_MRSIGNER;

DS:RCX.ISVEXTPRODID ← TMP\_SIG.ISVEXTPRODID;

DS:RCX.ISVPRODID ← TMP\_SIG.ISVPRODID;

DS:RCX.ISVSVN ← TMP\_SIG.ISVSVN;

DS:RCX.ISVFAMILYID ← TMP\_SIG.ISVFAMILYID;

DS:RCX.PADDING ← TMP\_SIG\_PADDING;

(\* Mark the SECS as initialized \*)

Update DS:RCX to initialized;

(\* Set RAX and ZF for success\*)

RFLAGS.ZF ← 0;

RAX ← 0;

EXIT:

RFLAGS.CF,PF,AF,OF,SF ← 0;

### Flags Affected

ZF is cleared if successful, otherwise ZF is set and RAX contains the error code. CF, PF, AF, OF, SF are cleared.

### Protected Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand is not properly aligned.</li> <li>If another instruction is modifying the SECS.</li> <li>If the enclave is already initialized.</li> <li>If the SECS.MRENCLAVE is in use.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If RCX does not resolve in an EPC page.</li> <li>If the memory address is not a valid, uninitialized SECS.</li> </ul>                      |

### 64-Bit Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand is not properly aligned.</li> <li>If another instruction is modifying the SECS.</li> <li>If the enclave is already initialized.</li> <li>If the SECS.MRENCLAVE is in use.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If RCX does not resolve in an EPC page.</li> <li>If the memory address is not a valid, uninitialized SECS.</li> </ul>                      |

## ELDB/ELDU/ELDBC/ELBUC—Load an EPC Page and Mark its State

| Opcode/<br>Instruction    | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|---------------------------|-------|------------------------------|--------------------------|--|
| EAX = 07H<br>ENCLS[ELDB]  | IR    | V/V                          | SGX1                     | This leaf function loads, verifies an EPC page and marks the page as blocked.                  |
| EAX = 08H<br>ENCLS[ELDU]  | IR    | V/V                          | SGX1                     | This leaf function loads, verifies an EPC page and marks the page as unblocked.                |
| EAX = 12H<br>ENCLS[ELDBC] | IR    | V/V                          | EAX[5]                   | This leaf function behaves like ELDB but with improved conflict handling for oversubscription. |
| EAX = 13H<br>ENCLS[ELBUC] | IR    | V/V                          | EAX[5]                   | This leaf function behaves like ELDU but with improved conflict handling for oversubscription. |

### Instruction Operand Encoding

| Op/En | EAX               |                            | RBX                             | RCX                             | RDX  |
|-------|-------------------|----------------------------|---------------------------------|---------------------------------|--|
| IR    | ELDB/ELDU<br>(In) | Return error<br>code (Out) | Address of the PAGEINFO<br>(In) | Address of the EPC page<br>(In) | Address of the version-<br>array slot (In) |

#### Description

This leaf function copies a page from regular main memory to the EPC. As part of the copying process, the page is cryptographically authenticated and decrypted. This instruction can only be executed when current privilege level is 0.

The ELDB leaf function sets the BLOCK bit in the EPCM entry for the destination page in the EPC after copying. The ELDU leaf function clears the BLOCK bit in the EPCM entry for the destination page in the EPC after copying.

RBX contains the effective address of a PAGEINFO structure; RCX contains the effective address of the destination EPC page; RDX holds the effective address of the version array slot that holds the version of the page.

The ELDBC/ELBUC leafs are very similar to ELDB and ELDU. They provide an error code on the concurrency conflict for any of the pages which need to acquire a lock. These include the destination, SECS, and VA slot.

The table below provides additional information on the memory parameter of ELDB/ELDU leaf functions.

### ELDB/ELDU/ELDBC/ELBUC Memory Parameter Semantics

| PAGEINFO                | PAGEINFO.SRCPGE         | PAGEINFO.PCMD           | PAGEINFO.SECS             | EPCPAGE                                | Version-Array Slot                     |
|-------------------------|-------------------------|-------------------------|---------------------------|--|--|
| Non-enclave read access | Non-enclave read access | Non-enclave read access | Enclave read/write access | Read/Write access permitted by Enclave | Read/Write access permitted by Enclave |

The error codes are:

**Table 40-28. ELDB/ELDU/ELDBC/ELBUC Return Value in RAX**

| Error Code (see Table 40-4) | Description             |
|-----------------------------|-------------------------|
| No Error                    | ELDB/ELDU successful.   |
| SGX_MAC_COMPARE_FAIL        | If the MAC check fails. |

## Concurrency Restrictions

Table 40-29. Base Concurrency Restrictions of ELDB/ELDU/ELDBC/ELBUC

| Leaf        | Parameter                  | Base Concurrency Restrictions |                       |                                    |
|-------------|----------------------------|-------------------------------|-----------------------|------------------------------------|
|             |                            | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| ELDB/ELDU/  | Target [DS:RCX]            | Exclusive                     | #GP                   | EPC_PAGE_CONFLICT_EXCEPTION        |
|             | VA [DS:RDX]                | Shared                        | #GP                   |                                    |
|             | SECS [DS:RBX]PAGEINFO.SECS | Shared                        | #GP                   |                                    |
| ELDBC/ELBUC | Target [DS:RCX]            | Exclusive                     | SGX_EPC_PAGE_CONFLICT | EPC_PAGE_CONFLICT_ERROR            |
|             | VA [DS:RDX]                | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|             | SECS [DS:RBX]PAGEINFO.SECS | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |

Table 40-30. Additional Concurrency Restrictions of ELDB/ELDU/ELDBC/ELBUC

| Leaf        | Parameter                  | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------------|----------------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|             |                            | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|             |                            | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| ELDB/ELDU/  | Target [DS:RCX]            | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | VA [DS:RDX]                | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | SECS [DS:RBX]PAGEINFO.SECS | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
| ELDBC/ELBUC | Target [DS:RCX]            | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | VA [DS:RDX]                | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | SECS [DS:RBX]PAGEINFO.SECS | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

## Operation

Temp Variables in ELDB/ELDU/ELDBC/ELBUC Operational Flow

| Name         | Type        | Size (Bits) | Description              |
|--------------|-------------|-------------|--------------------------|
| TMP_SRCPGE   | Memory page | 4KBytes     |                          |
| TMP_SECS     | Memory page | 4KBytes     |                          |
| TMP_PCMD     | PCMD        | 128 Bytes   |                          |
| TMP_HEADER   | MACHEADER   | 128 Bytes   |                          |
| TMP_VER      | UINT64      | 64          |                          |
| TMP_MAC      | UINT128     | 128         |                          |
| TMP_PK       | UINT128     | 128         | Page encryption/MAC key. |
| SCRATCH_PCMD | PCMD        | 128 Bytes   |                          |

(\* Check PAGEINFO and EPCPAGE alignment \*)

IF ( (DS:RBX is not 32Byte Aligned) or (DS:RCX is not 4KByte Aligned) )  
THEN #GP(0); FI;

## SGX INSTRUCTION REFERENCES

```
IF (DS:RCX does not resolve within an EPC)
    THEN #PF(DS:RCX); FI;
```

```
(* Check VASLOT alignment *)
IF (DS:RDX is not 8Byte aligned)
    THEN #GP(0); FI;
```

```
IF (DS:RDX does not resolve within an EPC)
    THEN #PF(DS:RDX); FI;
```

```
TMP_SRCPGE ← DS:RBX.SRCPGE;
TMP_SECS ← DS:RBX.SECONDS;
TMP_PCMD ← DS:RBX.PCMD;
```

```
(* Check alignment of PAGEINFO (RBX) linked parameters. Note: PCMD pointer is overlaid on top of PAGEINFO.SECINFO field *)
IF ( (DS:TMP_PCMD is not 128Byte aligned) or (DS:TMP_SRCPGE is not 4KByte aligned) )
    THEN #GP(0); FI;
```

```
(* Check concurrency of EPC by other Intel SGX instructions *)
IF (other instructions accessing EPC)
    THEN
        IF ((EAX==07h) OR (EAX==08h)) (* ELDB/ELDU *)
            THEN
                IF (<<VMX non-root operation>> AND
                    <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
                    THEN
                        VMCS.Exit_reason ← SGX_CONFLICT;
                        VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_EXCEPTION;
                        VMCS.Exit_qualification.error ← 0;
                        VMCS.Guest-physical_address ←
                            << translation of DS:RCX produced by paging >>;
                        VMCS.Guest-linear_address ← DS:RCX;
                        Deliver VMEXIT;
                    ELSE
                        #GP(0);
                FI;
            ELSE (* ELDBC/ELDUC *)
                IF (<<VMX non-root operation>> AND
                    <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
                    THEN
                        VMCS.Exit_reason ← SGX_CONFLICT;
                        VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_ERROR;
                        VMCS.Exit_qualification.error ← SGX_EPC_PAGE_CONFLICT;
                        VMCS.Guest-physical_address ←
                            << translation of DS:RCX produced by paging >>;
                        VMCS.Guest-linear_address ← DS:RCX;
                        Deliver VMEXIT;
                    ELSE
                        RFLAGS.ZF ← 1;
                        RFLAGS.CF ← 0;
                        RAX ← SGX_EPC_PAGE_CONFLICT;
                        GOTO ERROR_EXIT;
                FI;
```

```

    FI;
FI;

(* Check concurrency of EPC and VASLOT by other Intel SGX instructions *)
IF (Other instructions modifying VA slot)
    THEN
        IF ((EAX==07h) OR (EAX==08h)) (* ELDB/ELDU *)
            #GP(0);
            FI;
        ELSE (* ELDBC/ELDUC *)
            RFLAGS.ZF ← 1;
            RFLAGS.CF ← 0;
            RAX ← SGX_EPC_PAGE_CONFLICT;
            GOTO ERROR_EXIT;
FI;

(* Verify EPCM attributes of EPC page, VA, and SECS *)
IF (EPCM(DS:RCX).VALID = 1)
    THEN #PF(DS:RCX); FI;

IF ( (EPCM(DS:RDX & ~OFFFH).VALID = 0) or (EPCM(DS:RDX & ~OFFFH).PT ≠ PT_VA) )
    THEN #PF(DS:RDX); FI;

(* Copy PCMD into scratch buffer *)
SCRATCH_PCMD[1023: 0] ← DS:TMP_PCMD[1023:0];

(* Zero out TMP_HEADER*)
TMP_HEADER[sizeof(TMP_HEADER)-1: 0] ← 0;

TMP_HEADER.SECINFO ← SCRATCH_PCMD.SECINFO;
TMP_HEADER.RSVD ← SCRATCH_PCMD.RSVD;
TMP_HEADER.LINADDR ← DS:RBX.LINADDR;

(* Verify various attributes of SECS parameter *)
IF ( (TMP_HEADER.SECINFO.FLAGS.PT = PT_REG) or (TMP_HEADER.SECINFO.FLAGS.PT = PT_TCS) or
    (TMP_HEADER.SECINFO.FLAGS.PT = PT_TRIM) )
    THEN
        IF ( DS:TMP_SECS is not 4KByte aligned)
            THEN #GP(0) FI;
        IF (DS:TMP_SECS does not resolve within an EPC)
            THEN #PF(DS:TMP_SECS) FI;
        IF ( Other instructions modifying SECS)
            THEN
                IF ((EAX==07h) OR (EAX==08h)) (* ELDB/ELDU *)
                    #GP(0);
                    FI;
                ELSE (* ELDBC/ELDUC *)
                    RFLAGS.ZF ← 1;
                    RFLAGS.CF ← 0;
                    RAX ← SGX_EPC_PAGE_CONFLICT;
                    GOTO ERROR_EXIT;
FI;
FI;

```

## SGX INSTRUCTION REFERENCES

```
IF ( (TMP_HEADER.SECINFO.FLAGS.PT = PT_REG) or (TMP_HEADER.SECINFO.FLAGS.PT = PT_TCS) or
      (TMP_HEADER.SECINFO.FLAGS.PT = PT_TRIM) )
  THEN
    TMP_HEADER.EID ← DS:TMP_SECS.EID;
  ELSE
    (* These pages do not have any parent, and hence no EID binding *)
    TMP_HEADER.EID ← 0;
FI;

(* Copy 4KBytes SRCPGE to secure location *)
DS:RCX[32767: 0] ← DS:TMP_SRCPGE[32767: 0];
TMP_VER ← DS:RDX[63:0];

(* Decrypt and MAC page. AES_GCM_DEC has 2 outputs, {plain text, MAC} *)
(* Parameters for AES_GCM_DEC {Key, Counter, ..} *)
{DS:RCX, TMP_MAC} ← AES_GCM_DEC(CR_BASE_PK, TMP_VER << 32, TMP_HEADER, 128, DS:RCX, 4096);

IF ( (TMP_MAC ≠ DS:TMP_PCMD.MAC) )
  THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_MAC_COMPARE_FAIL;
    GOTO ERROR_EXIT;
FI;

(* Check version before committing *)
IF (DS:RDX ≠ 0)
  THEN #GP(0);
  ELSE
    DS:RDX ← TMP_VER;
FI;

(* Commit EPCM changes *)
EPCM(DS:RCX).PT ← TMP_HEADER.SECINFO.FLAGS.PT;
EPCM(DS:RCX).RWX ← TMP_HEADER.SECINFO.FLAGS.RWX;
EPCM(DS:RCX).PENDING ← TMP_HEADER.SECINFO.FLAGS.PENDING;
EPCM(DS:RCX).MODIFIED ← TMP_HEADER.SECINFO.FLAGS.MODIFIED;
EPCM(DS:RCX).PR ← TMP_HEADER.SECINFO.FLAGS.PR;
EPCM(DS:RCX).ENCLAVEADDRESS ← TMP_HEADER.LINADDR;

IF ( ((EAX = 07H) or (EAX = 12H)) and (TMP_HEADER.SECINFO.FLAGS.PT is NOT PT_SECS or PT_VA) )
  THEN
    EPCM(DS:RCX).BLOCKED ← 1;
  ELSE
    EPCM(DS:RCX).BLOCKED ← 0;
FI;

IF (TMP_HEADER.SECINFO.FLAGS.PT is PT_SECS)
  << store translation of DS:RCX produced by paging in SECS(DS:RCX).ENCLAVECONTEXT >>
FI;

EPCM(DS:RCX).VALID ← 1;

RAX ← 0;
RFLAGS.ZF ← 0;
```



ERROR\_EXIT:  
RFLAGS.CF,PF,AF,OF,SF ← 0;

### Flags Affected

Sets ZF if unsuccessful, otherwise cleared and RAX returns error code. Clears CF, PF, AF, OF, SF.

### Protected Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand effective address is outside the DS segment limit.</li> <li>If a memory operand is not properly aligned.</li> <li>If the instruction's EPC resource is in use by others.</li> <li>If the instruction fails to verify MAC.</li> <li>If the version-array slot is in use.</li> <li>If the parameters fail consistency checks.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand expected to be in EPC does not resolve to an EPC page.</li> <li>If one of the EPC memory operands has incorrect page type.</li> <li>If the destination EPC page is already valid.</li> </ul>   |

### 64-Bit Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand is non-canonical form.</li> <li>If a memory operand is not properly aligned.</li> <li>If the instruction's EPC resource is in use by others.</li> <li>If the instruction fails to verify MAC.</li> <li>If the version-array slot is in use.</li> <li>If the parameters fail consistency checks.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand expected to be in EPC does not resolve to an EPC page.</li> <li>If one of the EPC memory operands has incorrect page type.</li> <li>If the destination EPC page is already valid.</li> </ul>   |

## EMODPR—Restrict the Permissions of an EPC Page

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|----------------------------|-------|------------------------------|--------------------------|--|
| EAX = 0EH<br>ENCLS[EMODPR] | IR    | V/V                          | SGX2                     | This leaf function restricts the access rights associated with a EPC page in an initialized enclave. |

### Instruction Operand Encoding

| Op/En | EAX         |                         | RBX                       | RCX                                      |
|-------|-------------|-------------------------|---------------------------|--|
| IR    | EMODPR (In) | Return Error Code (Out) | Address of a SECINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function restricts the access rights associated with an EPC page in an initialized enclave. THE RWX bits of the SECINFO parameter are treated as a permissions mask; supplying a value that does not restrict the page permissions will have no effect. This instruction can only be executed when current privilege level is 0.

RBX contains the effective address of a SECINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of the EMODPR leaf function.

### EMODPR Memory Parameter Semantics

| SECINFO                              | EPCPAGE                                |
|--------------------------------------|--|
| Read access permitted by Non Enclave | Read/Write access permitted by Enclave |

The instruction faults if any of the following:

### EMODPR Faulting Conditions

|   |  |
|---|--|
| The operands are not properly aligned.    | If unsupported security attributes are set.                                      |
| The Enclave is not initialized.           | SECS is locked by another thread.  |
| The EPC page is locked by another thread. | RCX does not contain an effective address of an EPC page in the running enclave. |
| The EPC page is not valid.                |  |

The error codes are:

**Table 40-31. EMODPR Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EMODPR successful.  |
| SGX_PAGE_NOT_MODIFIABLE     | The EPC page cannot be modified because it is in the PENDING or MODIFIED state. |
| SGX_EPC_PAGE_CONFLICT       | Page is being written by EADD, EAUG, ECREATE, ELDU/B, EMODT, or EWB.            |

### Concurrency Restrictions

**Table 40-32. Base Concurrency Restrictions of EMODPR**

| Leaf   | Parameter       | Base Concurrency Restrictions |             |                                    |
|--------|-----------------|-------------------------------|-------------|------------------------------------|
|        |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EMODPR | Target [DS:RCX] | Shared                        | #GP         |                                    |

Table 40-33. Additional Concurrency Restrictions of EMODPR

| Leaf   | Parameter       | Additional Concurrency Restrictions             |                       |                          |             |                     |             |
|--------|-----------------|---|-----------------------|--------------------------|-------------|---------------------|-------------|
|        |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |                       | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |                 | Access  | On Conflict           | Access                   | On Conflict | Access              | On Conflict |
| EMODPR | Target [DS:RCX] | Exclusive                                       | SGX_EPC_PAGE_CONFLICT | Concurrent               |             | Concurrent          |             |

## Operation

### Temp Variables in EMODPR Operational Flow

| Name            | Type              | Size (bits) | Description  |
|-----------------|-------------------|-------------|--|
| TMP_SECS        | Effective Address | 32/64       | Physical address of SECS to which EPC operand belongs. |
| SCRATCH_SECINFO | SECINFO           | 512         | Scratch storage for holding the contents of DS:RBX.    |

IF (DS:RBX is not 64Byte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

SCRATCH\_SECINFO ← DS:RBX;

(\* Check for mis-configured SECINFO flags\*)

IF ( (SCRATCH\_SECINFO reserved fields are not zero ) or  
( (SCRATCH\_SECINFO.FLAGS.R is 0 and SCRATCH\_SECINFO.FLAGS.W is not 0 ) )  
THEN #GP(0); FI;

(\* Check concurrency with SGX1 or SGX2 instructions on the EPC page \*)

IF (SGX1 or other SGX2 instructions accessing EPC page)  
THEN #GP(0); FI;

IF (EPCM(DS:RCX).VALID is 0 )  
THEN #PF(DS:RCX); FI;

(\* Check the EPC page for concurrency \*)

IF (EPC page in use by another SGX2 instruction)  
THEN  
RFLAGS.ZF ← 1;  
RAX ← SGX\_EPC\_PAGE\_CONFLICT;  
GOTO DONE;

FI;

IF ( (EPCM(DS:RCX).PENDING is not 0 or (EPCM(DS:RCX).MODIFIED is not 0 ) )

THEN  
RFLAGS.ZF ← 1;  
RAX ← SGX\_PAGE\_NOT\_MODIFIABLE;

## SGX INSTRUCTION REFERENCES

```
GOTO DONE;
FI;

IF (EPCM(DS:RCX).PT is not PT_REG)
    THEN #PF(DS:RCX); FI;

TMP_SECS ← GET_SECS_ADDRESS

IF (TMP_SECS.ATTRIBUTES.INIT = 0)
    THEN #GP(0); FI;

(* Set the PR bit to indicate that permission restriction is in progress *)
EPCM(DS:RCX).PR ← 1;

(* Update EPCM permissions *)
EPCM(DS:RCX).R ← EPCM(DS:RCX).R & SCRATCH_SECINFO.FLAGS.R;
EPCM(DS:RCX).W ← EPCM(DS:RCX).W & SCRATCH_SECINFO.FLAGS.W;
EPCM(DS:RCX).X ← EPCM(DS:RCX).X & SCRATCH_SECINFO.FLAGS.X;

RFLAGS.ZF ← 0;
RAX ← 0;

DONE:
RFLAGS.CF,PF,AF,OF,SF ← 0;
```

### Flags Affected

Sets ZF if page is not modifiable or if other SGX2 instructions are executing concurrently, otherwise cleared. Clears CF, PF, AF, OF, SF.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the DS segment limit.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.  |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand is non-canonical form.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.                              |

## EMODT—Change the Type of an EPC Page

| Opcode/<br>Instruction    | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|---------------------------|-------|------------------------------|--------------------------|--|
| EAX = 0FH<br>ENCLS[EMODT] | IR    | V/V                          | SGX2                     | This leaf function changes the type of an existing EPC page. |

### Instruction Operand Encoding

| Op/En | EAX        |                         | RBX                       | RCX                                      |
|-------|------------|-------------------------|---------------------------|--|
| IR    | EMODT (In) | Return Error Code (Out) | Address of a SECINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function modifies the type of an EPC page. The security attributes are configured to prevent access to the EPC page at its new type until a corresponding invocation of the EACCEPT leaf confirms the modification. This instruction can only be executed when current privilege level is 0.

RBX contains the effective address of a SECINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of the EMODT leaf function.

### EMODT Memory Parameter Semantics

| SECINFO                              | EPCPAGE                                |
|--------------------------------------|--|
| Read access permitted by Non Enclave | Read/Write access permitted by Enclave |

The instruction faults if any of the following:

### EMODT Faulting Conditions

|   |  |
|---|--|
| The operands are not properly aligned.    | If unsupported security attributes are set.                                      |
| The Enclave is not initialized.           | SECS is locked by another thread.  |
| The EPC page is locked by another thread. | RCX does not contain an effective address of an EPC page in the running enclave. |
| The EPC page is not valid.                |  |

The error codes are:

**Table 40-34. EMODT Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EMODT successful.   |
| SGX_PAGE_NOT_MODIFIABLE     | The EPC page cannot be modified because it is in the PENDING or MODIFIED state. |
| SGX_EPC_PAGE_CONFLICT       | Page is being written by EADD, EAUG, ECREATE, ELDU/B, EMODPR, or EWB.           |

### Concurrency Restrictions

**Table 40-35. Base Concurrency Restrictions of EMODT**

| Leaf  | Parameter       | Base Concurrency Restrictions |                       |                                    |
|-------|-----------------|-------------------------------|-----------------------|------------------------------------|
|       |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| EMODT | Target [DS:RCX] | Exclusive                     | SGX_EPC_PAGE_CONFLICT | EPC_PAGE_CONFLICT_ERROR            |

**Table 40-36. Additional Concurrency Restrictions of EMODT**

| Leaf  | Parameter       | Additional Concurrency Restrictions             |                       |                          |             |                     |             |
|-------|-----------------|---|-----------------------|--------------------------|-------------|---------------------|-------------|
|       |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |                       | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|       |                 | Access  | On Conflict           | Access                   | On Conflict | Access              | On Conflict |
| EMODT | Target [DS:RCX] | Exclusive                                       | SGX_EPC_PAGE_CONFLICT | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EMODT Operational Flow**

| Name            | Type              | Size (bits) | Description  |
|-----------------|-------------------|-------------|--|
| TMP_SECS        | Effective Address | 32/64       | Physical address of SECS to which EPC operand belongs. |
| SCRATCH_SECINFO | SECINFO           | 512         | Scratch storage for holding the contents of DS:RBX.    |

IF (DS:RBX is not 64Byte Aligned)  
 THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
 THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

SCRATCH\_SECINFO ← DS:RBX;

(\* Check for mis-configured SECINFO flags\*)

IF ( (SCRATCH\_SECINFO reserved fields are not zero ) or  
 !(SCRATCH\_SECINFO.FLAGS.PT is PT\_TCS or SCRATCH\_SECINFO.FLAGS.PT is PT\_TRIM) )  
 THEN #GP(0); FI;

(\* Check concurrency with SGX1 instructions on the EPC page \*)

IF (other SGX1 instructions accessing EPC page)  
 THEN  
     RFLAGS.ZF ← 1;  
     RAX ← SGX\_EPC\_PAGE\_CONFLICT;  
     GOTO DONE;

FI;

IF (EPCM(DS:RCX).VALID is 0)  
 THEN #PF(DS:RCX); FI;

(\* Check the EPC page for concurrency \*)

IF (EPC page in use by another SGX2 instruction)  
 THEN  
     RFLAGS.ZF ← 1;  
     RAX ← SGX\_EPC\_PAGE\_CONFLICT;  
     GOTO DONE;

```

FI;

IF (!(EPCM(DS:RCX).PT is PT_REG or
      (EPCM(DS:RCX).PT is PT_TCS and SCRATCH_SECINFO.FLAGS.PT is PT_TRIM)))
    THEN #PF(DS:RCX); FI;

IF ( (EPCM(DS:RCX).PENDING is not 0 or (EPCM(DS:RCX).MODIFIED is not 0) )
    THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_PAGE_NOT_MODIFIABLE;
    GOTO DONE;
FI;

TMP_SECS ← GET_SECS_ADDRESS

IF (TMP_SECS.ATTRIBUTES.INIT = 0)
    THEN #GP(0); FI;

(* Update EPCM fields *)
EPCM(DS:RCX).PR ← 0;
EPCM(DS:RCX).MODIFIED ← 1;
EPCM(DS:RCX).R ← 0;
EPCM(DS:RCX).W ← 0;
EPCM(DS:RCX).X ← 0;
EPCM(DS:RCX).PT ← SCRATCH_SECINFO.FLAGS.PT;

RFLAGS.ZF ← 0;
RAX ← 0;

DONE:
RFLAGS.CF,PF,AF,OF,SF ← 0;

```

### Flags Affected

Sets ZF if page is not modifiable or if other SGX2 instructions are executing concurrently, otherwise cleared. Clears CF, PF, AF, OF, SF.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the DS segment limit.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.  |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand is non-canonical form.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked. |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.                              |

## EPA—Add Version Array

| Opcode/<br>Instruction  | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-------------------------|-------|------------------------------|--------------------------|---|
| EAX = 0AH<br>ENCLS[EPA] | IR    | V/V                          | SGX1                     | This leaf function adds a Version Array to the EPC. |

### Instruction Operand Encoding

| Op/En | EAX      | RBX                  | RCX                                    |
|-------|----------|----------------------|--|
| IR    | EPA (In) | PT_VA (In, Constant) | Effective address of the EPC page (In) |

### Description

This leaf function creates an empty version array in the EPC page whose logical address is given by DS:RCX, and sets up EPCM attributes for that page. At the time of execution of this instruction, the register RBX must be set to PT\_VA.

The table below provides additional information on the memory parameter of EPA leaf function.

### EPA Memory Parameter Semantics

|                                   |
|-----------------------------------|
| EPCPAGE                           |
| Write access permitted by Enclave |

### Concurrency Restrictions

Table 40-37. Base Concurrency Restrictions of EPA

| Leaf | Parameter   | Base Concurrency Restrictions |             |                                    |
|------|-------------|-------------------------------|-------------|------------------------------------|
|      |             | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EPA  | VA [DS:RCX] | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |

Table 40-38. Additional Concurrency Restrictions of EPA

| Leaf | Parameter   | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|------|-------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|      |             | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|      |             | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EPA  | VA [DS:RCX] | Concurrent                                      | L           | Concurrent               |             | Concurrent          |             |

### Operation

IF (RBX ≠ PT\_VA or DS:RCX is not 4KByte Aligned)  
THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
THEN #PF(DS:RCX); FI;

(\* Check concurrency with other Intel SGX instructions \*)

IF (Other Intel SGX instructions accessing the page)  
THEN

IF (<<VMX non-root operation>> AND <<ENABLE\_EPC\_VIRTUALIZATION\_EXTENSIONS>>)



```

THEN
    VMCS.Exit_reason ← SGX_CONFLICT;
    VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_EXCEPTION;
    VMCS.Exit_qualification.error ← 0;
    VMCS.Guest-physical_address ← <<< translation of DS:RCX produced by paging >>>;
    VMCS.Guest-linear_address ← DS:RCX;
    Deliver VMEXIT;
ELSE
    #GP(0);
FI;
FI;

```

(\* Check EPC page must be empty \*)

```

IF (EPCM(DS:RCX).VALID ≠ 0)
    THEN #PF(DS:RCX); FI;

```

(\* Clears EPC page \*)

```

DS:RCX[32767:0] ← 0;

```

```

EPCM(DS:RCX).PT ← PT_VA;
EPCM(DS:RCX).ENCLAVEADDRESS ← 0;
EPCM(DS:RCX).BLOCKED ← 0;
EPCM(DS:RCX).PENDING ← 0;
EPCM(DS:RCX).MODIFIED ← 0;
EPCM(DS:RCX).PR ← 0;
EPCM(DS:RCX).RWX ← 0;
EPCM(DS:RCX).VALID ← 1;

```

### Flags Affected

None

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand effective address is outside the DS segment limit.</li> <li>If a memory operand is not properly aligned.</li> <li>If another Intel SGX instruction is accessing the EPC page.</li> <li>If RBX is not set to PT_VA.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand is not an EPC page.</li> <li>If the EPC page is valid.</li> </ul>   |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand is non-canonical form.</li> <li>If a memory operand is not properly aligned.</li> <li>If another Intel SGX instruction is accessing the EPC page.</li> <li>If RBX is not set to PT_VA.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand is not an EPC page.</li> <li>If the EPC page is valid.</li> </ul>   |

## ERDINFO—Read Type and Status Information About an EPC Page

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 10H<br>ENCLS[ERDINFO] | IR    | V/V                          | EAX[6]                   | This leaf function returns type and status information about an EPC page. |

### Instruction Operand Encoding

| Op/En | EAX          | RBX                                | RCX                                      |
|-------|--------------|------------------------------------|--|
| IR    | ERDINFO (In) | Address of a RDINFO structure (In) | Address of the destination EPC page (In) |

### Description

This instruction reads type and status information about an EPC page and returns it in a RDINFO structure. The STATUS field of the structure describes the status of the page and determines the validity of the remaining fields. The FLAGS field returns the EPCM permissions of the page; the page type; and the BLOCKED, PENDING, MODIFIED, and PR status of the page. For enclave pages, the ENCLAVECONTEXT field of the structure returns the value of SECS.ENCLAVECONTEXT. For non-enclave pages (e.g., VA) ENCLAVECONTEXT returns 0.

For invalid or non-EPC pages, the instruction returns an information code indicating the page's status, in addition to populating the STATUS field.

ERDINFO returns an error code if the destination EPC page is being modified by a concurrent SGX instruction.

RBX contains the effective address of a RDINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of ERDINFO leaf function.

### ERDINFO Memory Parameter Semantics

| RDINFO                                     | EPCPAGE                          |
|--|----------------------------------|
| Read/Write access permitted by Non Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

### ERDINFO Faulting Conditions

|  |   |
|--|---|
| A memory operand effective address is outside the DS segment limit (32b mode). | A memory operand is not properly aligned.         |
| DS segment is unusable (32b mode).   | A page fault occurs in accessing memory operands. |
| A memory address is in a non-canonical form (64b mode).                        |   |

The error codes are:

**Table 40-39. ERDINFO Return Value in RAX**

| Error Code            | Value | Description   |
|-----------------------|-------|---|
| No Error              | 0     | ERDINFO successful.   |
| SGX_EPC_PAGE_CONFLICT |       | Failure due to concurrent operation of another SGX instruction. |
| SGX_PG_INVLD          |       | Target page is not a valid EPC page.                            |
| SGX_PG_NONEPC         |       | Page is not an EPC page.  |

Concurrency Restrictions

Table 40-40. Base Concurrency Restrictions of ERDINFO

| Leaf    | Parameter       | Base Concurrency Restrictions |                       |                                    |
|---------|-----------------|-------------------------------|-----------------------|------------------------------------|
|         |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| ERDINFO | Target [DS:RCX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |

Table 40-41. Additional Concurrency Restrictions of ERDINFO

| Leaf    | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| ERDINFO | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

Temp Variables in ERDINFO Operational Flow

| Name       | Type             | Size (Bits) | Description  |
|------------|------------------|-------------|--|
| TMP_SECS   | Physical Address | 64          | Physical address of the SECS of the page being modified. |
| TMP_RDINFO | Linear Address   | 64          | Address of the RDINFO structure.                         |

(\* check alignment of RDINFO structure (RBX) \*)  
 IF (DS:RBX is not 32Byte Aligned) THEN  
 #GP(0); FI;

(\* check alignment of the EPCPAGE (RCX) \*)  
 IF (DS:RCX is not 4KByte Aligned) THEN  
 #GP(0); FI;

(\* check that EPCPAGE (DS:RCX) is the address of an EPC page \*)  
 IF (DS:RCX does not resolve within EPC) THEN  
 RFLAGS.CF ← 1;  
 RFLAGS.ZF ← 0;  
 RAX ← SGX\_PG\_NONEPC;  
 goto DONE;  
 FI;

(\* Check the EPC page for concurrency \*)  
 IF (EPC page is being modified) THEN  
 RFLAGS.ZF = 1;  
 RFLAGS.CF = 0;  
 RAX = SGX\_EPC\_PAGE\_CONFLICT;  
 goto DONE;  
 FI;

(\* check page validity \*)  
 IF (EPCM(DS:RCX).VALID = 0) THEN  
 RFLAGS.CF = 1;

## SGX INSTRUCTION REFERENCES

```
RFLAGS.ZF = 0;
RAX = SGX_PG_INVLD;
goto DONE;
FI;

(* clear the fields of the RDINFO structure *)
TMP_RDINFO ← DS:RBX;
TMP_RDINFO.STATUS ← 0;
TMP_RDINFO.FLAGS ← 0;
TMP_RDINFO.ENCLAVECONTEXT ← 0;

(* store page info in RDINFO structure *)
TMP_RDINFO.FLAGS.RWX ← EPCM(DS:RCX).RWX;
TMP_RDINFO.FLAGS.PENDING ← EPCM(DS:RCX).PENDING;
TMP_RDINFO.FLAGS.MODIFIED ← EPCM(DS:RCX).MODIFIED;
TMP_RDINFO.FLAGS.PR ← EPCM(DS:RCX).PR;
TMP_RDINFO.FLAGS.PAGE_TYPE ← EPCM(DS:RCX).PAGE_TYPE;
TMP_RDINFO.FLAGS.BLOCKED ← EPCM(DS:RCX).BLOCKED;

(* read SECS.ENCLAVECONTEXT for enclave child pages *)
IF ((EPCM(DS:RCX).PAGE_TYPE = PT_REG) or
    (EPCM(DS:RCX).PAGE_TYPE = PT_TCS) or
    (EPCM(DS:RCX).PAGE_TYPE = PT_TRIM)
    ) THEN
    TMP_SECS ← Address of SECS for (DS:RCX);
    TMP_RDINFO.ENCLAVECONTEXT ← SECS(TMP_SECS).ENCLAVECONTEXT;
FI;

(* populate enclave information for SECS pages *)
IF (EPCM(DS:RCX).PAGE_TYPE = PT_SECS) THEN
    IF ((VMX non-root mode) and
        (ENABLE_EPC_VIRTUALIZATION_EXTENSIONS Execution Control = 1)
        ) THEN
        TMP_RDINFO.STATUS.CHILDPRESENT ←
            ((SECS(DS:RCX).CHLDCNT ≠ 0) or
             SECS(DS:RCX).VIRTCHILDCNT ≠ 0);
    ELSE
        TMP_RDINFO.STATUS.CHILDPRESENT ← (SECS(DS:RCX).CHLDCNT ≠ 0);
        TMP_RDINFO.STATUS.VIRTCHILDPRESENT ←
            (SECS(DS:RCX).VIRTCHILDCNT ≠ 0);
        TMP_RDINFO.ENCLAVECONTEXT ← SECS(DS:RCX).ENCLAVECONTEXT;
    FI;
FI;

RAX ← 0;
RFLAGS.ZF ← 0;
RFLAGS.CF ← 0;

DONE:
(* clear flags *)
RFLAGS.PF ← 0;
RFLAGS.AF ← 0;
RFLAGS.OF ← 0;
RFLAGS.SF ← 0;
```

**Flags Affected**

ZF is set if ERDINFO fails due to concurrent operation with another SGX instruction; otherwise cleared.

CF is set if page is not a valid EPC page or not an EPC page; otherwise cleared.

PF, AF, OF and SF are cleared.

**Protected Mode Exceptions**

- #GP(0)                    If a memory operand effective address is outside the DS segment limit.  
                          If DS segment is unusable.  
                          If a memory operand is not properly aligned.
- #PF(error code)        If a page fault occurs in accessing memory operands.

**64-Bit Mode Exceptions**

- #GP(0)                    If the memory address is in a non-canonical form.  
                          If a memory operand is not properly aligned.
- #PF(error code)        If a page fault occurs in accessing memory operands.

## EREMOVE—Remove a page from the EPC

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                     |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 03H<br>ENCLS[EREMOVE] | IR    | V/V                          | SGX1                     | This leaf function removes a page from the EPC. |

### Instruction Operand Encoding

| Op/En | EAX          | RCX                                    |
|-------|--------------|--|
| IR    | EREMOVE (In) | Effective address of the EPC page (In) |

#### Description

This leaf function causes an EPC page to be un-associated with its SECS and be marked as unused. This instruction leaf can only be executed when the current privilege level is 0.

The content of RCX is an effective address of an EPC page. The DS segment is used to create linear address. Segment override is not supported.

The instruction fails if the operand is not properly aligned or does not refer to an EPC page or the page is in use by another thread, or other threads are running in the enclave to which the page belongs. In addition the instruction fails if the operand refers to an SECS with associations.

### EREMOVE Memory Parameter Semantics

|                                   |
|-----------------------------------|
| EPCPAGE                           |
| Write access permitted by Enclave |

The instruction faults if any of the following:

### EREMOVE Faulting Conditions

|  |   |
|--|---|
| The memory operand is not properly aligned.              | The memory operand does not resolve in an EPC page.       |
| Refers to an invalid SECS.                               | Refers to an EPC page that is locked by another thread.   |
| Another Intel SGX instruction is accessing the EPC page. | RCX does not contain an effective address of an EPC page. |
| the EPC page refers to an SECS with associations.        |   |

The error codes are:

**Table 40-42. EREMOVE Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | EREMOVE successful.   |
| SGX_CHILD_PRESENT           | If the SECS still have enclave pages loaded into EPC.               |
| SGX_ENCLAVE_ACT             | If there are still logical processors executing inside the enclave. |

## Concurrency Restrictions

Table 40-43. Base Concurrency Restrictions of EREMOVE

| Leaf    | Parameter       | Base Concurrency Restrictions |             |                                    |
|---------|-----------------|-------------------------------|-------------|------------------------------------|
|         |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EREMOVE | Target [DS:RCX] | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |

Table 40-44. Additional Concurrency Restrictions of EREMOVE

| Leaf    | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EREMOVE | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

## Operation

## Temp Variables in EREMOVE Operational Flow

| Name     | Type              | Size (Bits) | Description                                     |
|----------|-------------------|-------------|---|
| TMP_SECS | Effective Address | 32/64       | Effective address of the SECS destination page. |

```
IF (DS:RCX is not 4KByte Aligned)
  THEN #GP(0); FI;
```

```
IF (DS:RCX does not resolve to an EPC page)
  THEN #PF(DS:RCX); FI;
```

```
TMP_SECS ← Get_SECS_ADDRESS();
```

```
(* Check the EPC page for concurrency *)
```

```
IF (EPC page being referenced by another Intel SGX instruction)
```

```
  THEN
```

```
    IF (<<VMX non-root operation>> AND <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
```

```
      THEN
```

```
        VMCS.Exit_reason ← SGX_CONFLICT;
```

```
        VMCS.Exit_qualification.code ← EPC_PAGE_CONFLICT_EXCEPTION;
```

```
        VMCS.Exit_qualification.error ← 0;
```

```
        VMCS.Guest-physical_address ← << translation of DS:RCX produced by paging >>;
```

```
        VMCS.Guest-linear_address ← DS:RCX;
```

```
        Deliver VMEXIT;
```

```
      ELSE
```

```
        #GP(0);
```

```
    FI;
```

```
FI;
```

```
(* if DS:RCX is already unused, nothing to do*)
```

```
IF ( (EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PT = PT_TRIM AND EPCM(DS:RCX).MODIFIED = 0))
```

```
  THEN GOTO DONE;
```

```
FI;
```

## SGX INSTRUCTION REFERENCES

```
IF ( (EPCM(DS:RCX).PT = PT_VA) OR
      ((EPCM(DS:RCX).PT = PT_TRIM) AND (EPCM(DS:RCX).MODIFIED = 0)) )
  THEN
    EPCM(DS:RCX).VALID ← 0;
    GOTO DONE;
FI;

IF (EPCM(DS:RCX).PT = PT_SECS)
  THEN
    IF (DS:RCX has an EPC page associated with it)
      THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_CHILD_PRESENT;
        GOTO ERROR_EXIT;

    FI;
    (* treat SECS as having a child page when VIRTCHILDCNT is non-zero *)
    IF (<<in VMX non-root operation>> AND
        <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>> AND
        (SECS(DS:RCX).VIRTCHILDCNT ≠ 0))
      THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_CHILD_PRESENT
        GOTO ERROR_EXIT

    FI;
    EPCM(DS:RCX).VALID ← 0;
    GOTO DONE;
FI;

IF (Other threads active using SECS)
  THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_ENCLAVE_ACT;
    GOTO ERROR_EXIT;
FI;

IF ( (EPCM(DS:RCX).PT is PT_REG) or (EPCM(DS:RCX).PT is PT_TCS) or (EPCM(DS:RCX).PT is PT_TRIM) )
  THEN
    EPCM(DS:RCX).VALID ← 0;
    GOTO DONE;
FI;

DONE:
RAX ← 0;
RFLAGS.ZF ← 0;

ERROR_EXIT:
RFLAGS.CF,PF,AF,OF,SF ← 0;
```

### Flags Affected

Sets ZF if unsuccessful, otherwise cleared and RAX returns error code. Clears CF, PF, AF, OF, SF.



**Protected Mode Exceptions**

- #GP(0)                If a memory operand effective address is outside the DS segment limit.  
                          If a memory operand is not properly aligned.  
                          If another Intel SGX instruction is accessing the page.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                          If the memory operand is not an EPC page.

**64-Bit Mode Exceptions**

- #GP(0)                If the memory operand is non-canonical form.  
                          If a memory operand is not properly aligned.  
                          If another Intel SGX instruction is accessing the page.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                          If the memory operand is not an EPC page.

## ETRAK—Activates EBLOCK Checks

| Opcode/<br>Instruction    | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                 |
|---------------------------|-------|------------------------------|--------------------------|---|
| EAX = 0CH<br>ENCLS[ETRAK] | IR    | V/V                          | SGX1                     | This leaf function activates EBLOCK checks. |

### Instruction Operand Encoding

| Op/En | EAX        |                         | RCX                                      |
|-------|------------|-------------------------|--|
| IR    | ETRAK (In) | Return error code (Out) | Pointer to the SECS of the EPC page (In) |

### Description

This leaf function provides the mechanism for hardware to track that software has completed the required TLB address clears successfully. The instruction can only be executed when the current privilege level is 0.

The content of RCX is an effective address of an EPC page.

The table below provides additional information on the memory parameter of ETRAK leaf function.

### ETRAK Memory Parameter Semantics

|  |
|--|
| EPCPAGE                                |
| Read/Write access permitted by Enclave |

The error codes are:

**Table 40-45. ETRAK Return Value in RAX**

| Error Code (see Table 40-4) | Description   |
|-----------------------------|---|
| No Error                    | ETRAK successful.   |
| SGX_PREV_TRK_INCMPL         | All processors did not complete the previous shoot-down sequence. |

### Concurrency Restrictions

**Table 40-46. Base Concurrency Restrictions of ETRAK**

| Leaf  | Parameter     | Base Concurrency Restrictions |             |                                    |
|-------|---------------|-------------------------------|-------------|------------------------------------|
|       |               | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| ETRAK | SECS [DS:RCX] | Shared                        | #GP         |                                    |

**Table 40-47. Additional Concurrency Restrictions of ETRAK**

| Leaf  | Parameter     | Additional Concurrency Restrictions             |             |                          |             |                   |                       |
|-------|---------------|---|-------------|--------------------------|-------------|-------------------|-----------------------|
|       |               | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRAK, ETRAKC |                       |
|       |               | Access  | On Conflict | Access                   | On Conflict | Access            | On Conflict           |
| ETRAK | SECS [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Exclusive         | SGX_EPC_PAGE_CONFLICT |

**Operation**

```
IF (DS:RCX is not 4KByte Aligned)
  THEN #GP(0); FI;
```

```
IF (DS:RCX does not resolve within an EPC)
  THEN #PF(DS:RCX); FI;
```

(\* Check concurrency with other Intel SGX instructions \*)

```
IF (Other Intel SGX instructions using tracking facility on this SECS)
  THEN
    IF (<<VMX non-root operation>> AND <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
      THEN
        VMCS.Exit_reason ← SGX_CONFLICT;
        VMCS.Exit_qualification.code ← TRACKING_RESOURCE_CONFLICT;
        VMCS.Exit_qualification.error ← 0;
        VMCS.Guest-physical_address ← SECS(TMP_SECS).ENCLAVECONTEXT;
        VMCS.Guest-linear_address ← 0;
        Deliver VMEXIT;
      ELSE
        #GP(0);
    FI;
  FI;
```

```
FI;
```

```
IF (EPCM(DS:RCX).VALID = 0)
  THEN #PF(DS:RCX); FI;
```

```
IF (EPCM(DS:RCX).PT ≠ PT_SECS)
  THEN #PF(DS:RCX); FI;
```

(\* All processors must have completed the previous tracking cycle\*)

```
IF ( (DS:RCX).TRACKING ≠ 0 )
  THEN
    IF (<<VMX non-root operation>> AND <<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>>)
      THEN
        VMCS.Exit_reason ← SGX_CONFLICT;
        VMCS.Exit_qualification.code ← TRACKING_REFERENCE_CONFLICT;
        VMCS.Exit_qualification.error ← 0;
        VMCS.Guest-physical_address ← SECS(TMP_SECS).ENCLAVECONTEXT;
        VMCS.Guest-linear_address ← 0;
        Deliver VMEXIT;
      FI;
    RFLAGS.ZF ← 1;
    RAX ← SGX_PREV_TRK_INCMPL;
    GOTO DONE;
  ELSE
    RAX ← 0;
    RFLAGS.ZF ← 0;
  FI;
```

```
DONE:
RFLAGS.CF,PF,AF,OF,SF ← 0;
```

**Flags Affected**

Sets ZF if SECS is in use or invalid, otherwise cleared. Clears CF, PF, AF, OF, SF.

**Protected Mode Exceptions**

- #GP(0)                If a memory operand effective address is outside the DS segment limit.  
                          If a memory operand is not properly aligned.  
                          If another thread is concurrently using the tracking facility on this SECS.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                          If a memory operand is not an EPC page.

**64-Bit Mode Exceptions**

- #GP(0)                If a memory operand is non-canonical form.  
                          If a memory operand is not properly aligned.  
                          If the specified EPC resource is in use.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                          If a memory operand is not an EPC page.

## ETRACKC—Activates EBLOCK Checks

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                 |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 11H<br>ENCLS[ETRACKC] | IR    | V/V                          | EAX[6]                   | This leaf function activates EBLOCK checks. |

### Instruction Operand Encoding

| Op/En | EAX            |                         | RCX   |                                   |
|-------|----------------|-------------------------|---|-----------------------------------|
| IR    | ETRACK<br>(In) | Return error code (Out) | Address of the destination EPC page<br>(In, EA) | Address of the SECS page (In, EA) |

### Description

The ETRACKC instruction is thread safe variant of ETRACK leaf and can be executed concurrently with other CPU threads operating on the same SECS.

This leaf function provides the mechanism for hardware to track that software has completed the required TLB address clears successfully. The instruction can only be executed when the current privilege level is 0.

The content of RCX is an effective address of an EPC page.

The table below provides additional information on the memory parameter of ETRACK leaf function.

### ETRACKC Memory Parameter Semantics

|  |
|--|
| EPCPAGE                                |
| Read/Write access permitted by Enclave |

The error codes are:

**Table 40-48. ETRACKC Return Value in RAX**

| Error Code             | Value | Description   |
|------------------------|-------|---|
| No Error               | 0     | ETRACKC successful.   |
| SGX_EPC_PAGE_CONFLICT  | 7     | Failure due to concurrent operation of another SGX instruction. |
| SGX_PG_INVLD           | 6     | Target page is not a VALID EPC page.                            |
| SGX_PREV_TRK_INCMPL    | 17    | All processors did not complete the previous tracking sequence. |
| SGX_TRACK_NOT_REQUIRED | 27    | Target page type does not require tracking.                     |

### Concurrency Restrictions

**Table 40-49. Base Concurrency Restrictions of ETRACKC**

| Leaf    | Parameter       | Base Concurrency Restrictions |                       |                                    |
|---------|-----------------|-------------------------------|-----------------------|------------------------------------|
|         |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| ETRACKC | Target [DS:RCX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|         | SECS implicit   | Concurrent                    |                       |                                    |

**Table 40-50. Additional Concurrency Restrictions of ETRACKC**

| Leaf    | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |                       |
|---------|-----------------|---|-------------|--------------------------|-------------|---------------------|-----------------------|
|         |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |                       |
|         |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict           |
| ETRACKC | Target [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |                       |
|         | SECS implicit   | Concurrent                                      |             | Concurrent               |             | Exclusive           | SGX_EPC_PAGE_CONFLICT |

**Operation**

**Temp Variables in ETRACKC Operational Flow**

| Name     | Type             | Size (Bits) | Description  |
|----------|------------------|-------------|--|
| TMP_SECS | Physical Address | 64          | Physical address of the SECS of the page being modified. |

(\* check alignment of EPCPAGE (RCX) \*)  
 IF (DS:RCX is not 4KByte Aligned) THEN  
 #GP(0); FI;

(\* check that EPCPAGE (DS:RCX) is the address of an EPC page \*)  
 IF (DS:RCX does not resolve within an EPC) THEN  
 #PF(DS:RCX, PFEC.SGX); FI;

(\* Check the EPC page for concurrency \*)  
 IF (EPC page is being modified) THEN  
 RFLAGS.ZF ← 1;  
 RFLAGS.CF ← 0;  
 RAX ← SGX\_EPC\_PAGE\_CONFLICT;  
 goto DONE\_POST\_LOCK\_RELEASE;  
 FI;

(\* check to make sure the page is valid \*)  
 IF (EPCM(DS:RCX).VALID = 0) THEN  
 RFLAGS.ZF ← 1;  
 RFLAGS.CF ← 0;  
 RAX ← SGX\_PG\_INVLD;  
 GOTO DONE;  
 FI;

(\* find out the target SECS page \*)  
 IF (EPCM(DS:RCX).PT is PT\_REG or PT\_TCS or PT\_TRIM) THEN  
 TMP\_SECS ← Obtain SECS through EPCM(DS:RCX).ENCLAVESECS;  
 ELSE IF (EPCM(DS:RCX).PT is PT\_SECS) THEN  
 TMP\_SECS ← Obtain SECS through (DS:RCX);  
 ELSE  
 RFLAGS.ZF ← 0;  
 RFLAGS.CF ← 1;  
 RAX ← SGX\_TRACK\_NOT\_REQUIRED;  
 GOTO DONE;  
 FI;

```
(* Check concurrency with other Intel SGX instructions *)
IF (Other Intel SGX instructions using tracking facility on this SECS) THEN
  IF ((VMX non-root mode) and
    (ENABLE_EPC_VIRTUALIZATION_EXTENSIONS Execution Control = 1)) THEN
    VMCS.Exit_reason ← SGX_CONFLICT;
    VMCS.Exit_qualification.code ← TRACKING_RESOURCE_CONFLICT;
    VMCS.Exit_qualification.error ← 0;
    VMCS.Guest-physical_address ←
      SECS(TMP_SECS).ENCLAVECONTEXT;
    VMCS.Guest-linear_address ← 0;
    Deliver VMEXIT;
  FI;
```

```
RFLAGS.ZF ← 1;
RFLAGS.CF ← 0;
RAX ← SGX_EPC_PAGE_CONFLICT;
GOTO DONE;
```

```
FI;

(* All processors must have completed the previous tracking cycle*)
IF ((TMP_SECS).TRACKING ≠ 0)
THEN
  IF ((VMX non-root mode) and
    (ENABLE_EPC_VIRTUALIZATION_EXTENSIONS Execution Control = 1)) THEN
    VMCS.Exit_reason ← SGX_CONFLICT;
    VMCS.Exit_qualification.code ← TRACKING_REFERENCE_CONFLICT;
    VMCS.Exit_qualification.error ← 0;
    VMCS.Guest-physical_address ←
      SECS(TMP_SECS).ENCLAVECONTEXT;
    VMCS.Guest-linear_address ← 0;
    Deliver VMEXIT;
  FI;
```

```
RFLAGS.ZF ← 1;
RFLAGS.CF ← 0;
RAX ← SGX_PREV_TRK_INCMPL;
GOTO DONE;
```

```
FI;
```

```
RFLAGS.ZF ← 0;
RFLAGS.CF ← 0;
RAX ← 0;
```

```
DONE:
(* clear flags *)
RFLAGS.PF,AF,OF,SF ← 0;
```

### Flags Affected

ZF is set if ETRACKC fails due to concurrent operations with another SGX instructions or target page is an invalid EPC page or tracking is not completed on SECS page; otherwise cleared.

CF is set if target page is not of a type that requires tracking; otherwise cleared.

PF, AF, OF and SF are cleared.

**Protected Mode Exceptions**

- #GP(0) If the memory operand violates access-control policies of DS segment.  
If DS segment is unusable.
- #PF(error code) If the memory operand is not properly aligned.  
If the memory operand expected to be in EPC does not resolve to an EPC page.  
If a page fault occurs in access memory operand.

**64-Bit Mode Exceptions**

- #GP(0) If a memory address is in a non-canonical form.  
If a memory operand is not properly aligned.
- #PF(error code) If the memory operand expected to be in EPC does not resolve to an EPC page.  
If a page fault occurs in access memory operand.



## EWB—Invalidate an EPC Page and Write out to Main Memory

| Opcode/<br>Instruction  | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|-------------------------|-------|------------------------------|--------------------------|--|
| EAX = 0BH<br>ENCLS[EWB] | IR    | V/V                          | SGX1                     | This leaf function invalidates an EPC page and writes it out to main memory. |

### Instruction Operand Encoding

| Op/En | EAX      |                  | RBX                         | RCX                          | RDX                       |
|-------|----------|------------------|-----------------------------|------------------------------|---------------------------|
| IR    | EWB (In) | Error code (Out) | Address of an PAGEINFO (In) | Address of the EPC page (In) | Address of a VA slot (In) |

### Description

This leaf function copies a page from the EPC to regular main memory. As part of the copying process, the page is cryptographically protected. This instruction can only be executed when current privilege level is 0.

The table below provides additional information on the memory parameter of EPA leaf function.

### EWB Memory Parameter Semantics

| PAGEINFO           | PAGEINFO.SRCPGE    | PAGEINFO.PCMD      | EPCPAGE        | VASLOT         |
|--------------------|--------------------|--------------------|----------------|----------------|
| Non-EPC R/W access | Non-EPC R/W access | Non-EPC R/W access | EPC R/W access | EPC R/W access |

The error codes are:

**Table 40-51. EWB Return Value in RAX**

| Error Code (see Table 40-4) | Description  |
|-----------------------------|--|
| No Error                    | EWB successful.  |
| SGX_PAGE_NOT_BLOCKED        | If page is not marked as blocked.                        |
| SGX_NOT_TRACKED             | If EWB is racing with ETRACK instruction.                |
| SGX_VA_SLOT_OCCUPIED        | Version array slot contained valid entry.                |
| SGX_CHILD_PRESENT           | Child page present while attempting to page out enclave. |

### Concurrency Restrictions

**Table 40-52. Base Concurrency Restrictions of EWB**

| Leaf | Parameter       | Base Concurrency Restrictions |             |                                    |
|------|-----------------|-------------------------------|-------------|------------------------------------|
|      |                 | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EWB  | Source [DS:RCX] | Exclusive                     | #GP         | EPC_PAGE_CONFLICT_EXCEPTION        |
|      | VA [DS:RDX]     | Shared                        | #GP         |                                    |

**Table 40-53. Additional Concurrency Restrictions of EWB**

| Leaf | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|      |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|      |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EWB  | Source [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|      | VA [DS:RDX]     | Concurrent                                      |             | Concurrent               |             | Exclusive           |             |

Operation

Temp Variables in EWB Operational Flow

| Name               | Type        | Size (Bytes) | Description |
|--------------------|-------------|--------------|-------------|
| TMP_SRCPGE         | Memory page | 4096         |             |
| TMP_PCMD           | PCMD        | 128          |             |
| TMP_SECS           | SECS        | 4096         |             |
| TMP_BPEPOCH        | UINT64      | 8            |             |
| TMP_BPREFCOUNT     | UINT64      | 8            |             |
| TMP_HEADER         | MAC Header  | 128          |             |
| TMP_PCMD_ENCLAVEID | UINT64      | 8            |             |
| TMP_VER            | UINT64      | 8            |             |
| TMP_PK             | UINT128     | 16           |             |

IF ( (DS:RBX is not 32Byte Aligned) or (DS:RCX is not 4KByte Aligned) )  
 THEN #GP(0); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

IF (DS:RDX is not 8Byte Aligned)  
 THEN #GP(0); FI;

IF (DS:RDX does not resolve within an EPC)  
 THEN #PF(DS:RDX); FI;

(\* EPCPAGE and VASLOT should not resolve to the same EPC page\*)  
 IF (DS:RCX and DS:RDX resolve to the same EPC page)  
 THEN #GP(0); FI;

TMP\_SRCPGE ← DS:RBX.SRCPGE;  
 (\* Note PAGEINFO.PCMD is overlaid on top of PAGEINFO.SECINFO \*)  
 TMP\_PCMD ← DS:RBX.PCMD;

If (DS:RBX.LINADDR ≠ 0) OR (DS:RBX.SECS ≠ 0)  
 THEN #GP(0); FI;

IF ( (DS:TMP\_PCMD is not 128Byte Aligned) or (DSTMP\_SRCPGE is not 4KByte Aligned) )  
 THEN #GP(0); FI;

(\* Check for concurrent Intel SGX instruction access to the page \*)  
 IF (Other Intel SGX instruction is accessing page)  
 THEN  
     IF (<<VMX non-root operation>> AND <<ENABLE\_EPC\_VIRTUALIZATION\_EXTENSIONS>>)  
     THEN  
         VMCS.Exit\_reason ← SGX\_CONFLICT;  
         VMCS.Exit\_qualification.code ← EPC\_PAGE\_CONFLICT\_EXCEPTION;  
         VMCS.Exit\_qualification.error ← 0;  
         VMCS.Guest-physical\_address ← << translation of DS:RCX produced by paging >>;

```

        VMCS.Guest-linear_address ← DS:RCX;
        Deliver VMEXIT;
        ELSE
            #GP(0);
    FI;
FI;

(*Check if the VA Page is being removed or changed*)
IF (VA Page is being modified)
    THEN #GP(0); FI;

(* Verify that EPCPAGE and VASLOT page are valid EPC pages and DS:RDX is VA *)
IF (EPCM(DS:RCX).VALID = 0)
    THEN #PF(DS:RCX); FI;

IF ( (EPCM(DS:RDX & ~OFFFH).VALID = 0) or (EPCM(DS:RDX & ~FFFH).PT is not PT_VA) )
    THEN #PF(DS:RDX); FI;

(* Perform page-type-specific exception checks *)
IF ( (EPCM(DS:RCX).PT is PT_REG) or (EPCM(DS:RCX).PT is PT_TCS) or (EPCM(DS:RCX).PT is PT_TRIM) )
    THEN
        TMP_SECS = Obtain SECS through EPCM(DS:RCX)
        (* Check that EBLOCK has occurred correctly *)
        IF (EBLOCK is not correct)
            THEN #GP(0); FI;
FI;

RFLAGS.ZF,CF,PF,AF,OF,SF ← 0;
RAX ← 0;

(* Perform page-type-specific checks *)
IF ( (EPCM(DS:RCX).PT is PT_REG) or (EPCM(DS:RCX).PT is PT_TCS) or (EPCM(DS:RCX).PT is PT_TRIM) )
    THEN
        (* check to see if the page is evictable *)
        IF (EPCM(DS:RCX).BLOCKED = 0)
            THEN
                RAX ← SGX_PAGE NOT_BLOCKED;
                RFLAGS.ZF ← 1;
                GOTO ERROR_EXIT;
        FI;
        (* Check if tracking done correctly *)
        IF (Tracking not correct)
            THEN
                RAX ← SGX_NOT_TRACKED;
                RFLAGS.ZF ← 1;
                GOTO ERROR_EXIT;
        FI;

        (* Obtain EID to establish cryptographic binding between the paged-out page and the enclave *)
        TMP_HEADER.EID ← TMP_SECS.EID;

        (* Obtain EID as an enclave handle for software *)
        TMP_PCMD_ENCLAVEID ← TMP_SECS.EID;
    ELSE IF (EPCM(DS:RCX).PT is PT_SECS)

```

```

(*check that there are no child pages inside the enclave *)
IF (DS:RCX has an EPC page associated with it)
    THEN
        RAX ← SGX_CHILD_PRESENT;
        RFLAGS.ZF ← 1;
        GOTO ERROR_EXIT;
FI;
(* treat SECS as having a child page when VIRTCHILDCNT is non-zero *)
IF (<<in VMX non-root operation>> AND
<<ENABLE_EPC_VIRTUALIZATION_EXTENSIONS>> AND
(SECS(DS:RCX).VIRTCHILDCNT ≠ 0))
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_CHILD_PRESENT;
        GOTO ERROR_EXIT;
FI;
TMP_HEADER.EID ← 0;
(* Obtain EID as an enclave handle for software *)
TMP_PCMD_ENCLAVEID ← (DS:RCX).EID;
ELSE IF (EPCM(DS:RCX).PT is PT_VA)
    TMP_HEADER.EID ← 0; // Zero is not a special value
    (* No enclave handle for VA pages*)
    TMP_PCMD_ENCLAVEID ← 0;
FI;

(* Zero out TMP_HEADER*)
TMP_HEADER[ sizeof(TMP_HEADER)-1 : 0] ← 0;

TMP_HEADER.LINADDR ← EPCM(DS:RCX).ENCLAVEADDRESS;
TMP_HEADER.SECINFO.FLAGS.PT ← EPCM(DS:RCX).PT;
TMP_HEADER.SECINFO.FLAGS.RWX ← EPCM(DS:RCX).RWX;
TMP_HEADER.SECINFO.FLAGS.PENDING ← EPCM(DS:RCX).PENDING;
TMP_HEADER.SECINFO.FLAGS.MODIFIED ← EPCM(DS:RCX).MODIFIED;
TMP_HEADER.SECINFO.FLAGS.PR ← EPCM(DS:RCX).PR;

(* Encrypt the page, DS:RCX could be encrypted in place. AES-GCM produces 2 values, {ciphertext, MAC}. *)
(* AES-GCM input parameters: key, GCM Counter, MAC_HDR, MAC_HDR_SIZE, SRC, SRC_SIZE*)
{DS:TMP_SRCPGE, DS:TMP_PCMD.MAC} ← AES_GCM_ENC(CR_BASE_PK), (TMP_VER << 32),
    TMP_HEADER, 128, DS:RCX, 4096);

(* Write the output *)
Zero out DS:TMP_PCMD.SECINFO
DS:TMP_PCMD.SECINFO.FLAGS.PT ← EPCM(DS:RCX).PT;
DS:TMP_PCMD.SECINFO.FLAGS.RWX ← EPCM(DS:RCX).RWX;
DS:TMP_PCMD.SECINFO.FLAGS.PENDING ← EPCM(DS:RCX).PENDING;
DS:TMP_PCMD.SECINFO.FLAGS.MODIFIED ← EPCM(DS:RCX).MODIFIED;
DS:TMP_PCMD.SECINFO.FLAGS.PR ← EPCM(DS:RCX).PR;
DS:TMP_PCMD.RESERVED ← 0;
DS:TMP_PCMD.ENCLAVEID ← TMP_PCMD_ENCLAVEID;
DS:RBX.LINADDR ← EPCM(DS:RCX).ENCLAVEADDRESS;

(*Check if version array slot was empty *)
IF ([DS.RDX])
    THEN

```

```
RAX ← SGX_VA_SLOT_OCCUPIED
RFLAGS.CF ← 1;
```

```
FI;
```

```
(* Write version to Version Array slot *)
[DS.RDX] ← TMP_VER;
```

```
(* Free up EPCM Entry *)
EPCM.(DS:RCX).VALID ← 0;
ERROR_EXIT:
```

### Flags Affected

ZF is set if page is not blocked, not tracked, or a child is present. Otherwise cleared.

CF is set if VA slot is previously occupied, Otherwise cleared.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand effective address is outside the DS segment limit.</li> <li>If a memory operand is not properly aligned.</li> <li>If the EPC page and VASLOT resolve to the same EPC page.</li> <li>If another Intel SGX instruction is concurrently accessing either the target EPC, VA, or SECS pages.</li> <li>If the tracking resource is in use.</li> <li>If the EPC page or the version array page is invalid.</li> <li>If the parameters fail consistency checks.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand is not an EPC page.</li> <li>If one of the EPC memory operands has incorrect page type.</li> </ul>  |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <ul style="list-style-type: none"> <li>If a memory operand is non-canonical form.</li> <li>If a memory operand is not properly aligned.</li> <li>If the EPC page and VASLOT resolve to the same EPC page.</li> <li>If another Intel SGX instruction is concurrently accessing either the target EPC, VA, or SECS pages.</li> <li>If the tracking resource is in use.</li> <li>If the EPC page or the version array page in invalid.</li> <li>If the parameters fail consistency checks.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory operands.</li> <li>If a memory operand is not an EPC page.</li> <li>If one of the EPC memory operands has incorrect page type.</li> </ul>  |

## 40.4 INTEL® SGX USER LEAF FUNCTION REFERENCE

Leaf functions available with the ENCLU instruction mnemonic are covered in this section. In general, each instruction leaf requires EAX to specify the leaf function index and/or additional registers specifying leaf-specific input parameters. An instruction operand encoding table provides details of the implicitly-encoded register usage and associated input/output semantics.

In many cases, an input parameter specifies an effective address associated with a memory object inside or outside the EPC, the memory addressing semantics of these memory objects are also summarized in a separate table.

## EACCEPT—Accept Changes to an EPC Page

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 05H<br>ENCLU[EACCEPT] | IR    | V/V                          | SGX2                     | This leaf function accepts changes made by system software to an EPC page in the running enclave. |

### Instruction Operand Encoding

| Op/En | EAX          |                         | RBX                       | RCX                                      |
|-------|--------------|-------------------------|---------------------------|--|
| IR    | EACCEPT (In) | Return Error Code (Out) | Address of a SECINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function accepts changes to a page in the running enclave by verifying that the security attributes specified in the SECINFO match the security attributes of the page in the EPCM. This instruction leaf can only be executed when inside the enclave.

RBX contains the effective address of a SECINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of the EACCEPT leaf function.

### EACCEPT Memory Parameter Semantics

| SECINFO                              | EPCPAGE (Destination)            |
|--------------------------------------|----------------------------------|
| Read access permitted by Non Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

### EACCEPT Faulting Conditions

|  |  |
|--|--|
| The operands are not properly aligned.                                 | RBX does not contain an effective address in an EPC page in the running enclave. |
| The EPC page is locked by another thread.                              | RCX does not contain an effective address of an EPC page in the running enclave. |
| The EPC page is not valid.   | Page type is PT_REG and MODIFIED bit is 0.                                       |
| SECINFO contains an invalid request.                                   | Page type is PT_TCS or PT_TRIM and PENDING bit is 0 and MODIFIED bit is 1.       |
| If security attributes of the SECINFO page make the page inaccessible. |  |

The error codes are:

**Table 40-54. EACCEPT Return Value in RAX**

| Error Code (see Table 40-4)  | Description   |
|------------------------------|---|
| No Error                     | EACCEPT successful.   |
| SGX_PAGE_ATTRIBUTES_MISMATCH | The attributes of the target EPC page do not match the expected values. |
| SGX_NOT_TRACKED              | The OS did not complete an ETRACK on the target page.                   |

Concurrency Restrictions

**Table 40-55. Base Concurrency Restrictions of EACCEPT**

| Leaf    | Parameter        | Base Concurrency Restrictions |             |                                    |
|---------|------------------|-------------------------------|-------------|------------------------------------|
|         |                  | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EACCEPT | Target [DS:RCX]  | Shared                        | #GP         |                                    |
|         | SECINFO [DS:RBX] | Concurrent                    |             |                                    |

**Table 40-56. Additional Concurrency Restrictions of EACCEPT**

| Leaf    | Parameter        | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                  | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                  | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EACCEPT | Target [DS:RCX]  | Exclusive                                       | #GP         | Concurrent               |             | Concurrent          |             |
|         | SECINFO [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

**Temp Variables in EACCEPT Operational Flow**

| Name            | Type              | Size (bits) | Description   |
|-----------------|-------------------|-------------|---|
| TMP_SECS        | Effective Address | 32/64       | Physical address of SECS to which EPC operands belongs. |
| SCRATCH_SECINFO | SECINFO           | 512         | Scratch storage for holding the contents of DS:RBX.     |

IF (DS:RBX is not 64Byte Aligned)  
 THEN #GP(0); FI;

IF (DS:RBX is not within CR\_ELRANGE)  
 THEN #GP(0); FI;

IF (DS:RBX does not resolve within an EPC)  
 THEN #PF(DS:RBX); FI;

IF ( (EPCM(DS:RBX &~FFFH).VALID = 0) or (EPCM(DS:RBX &~FFFH).R = 0) or (EPCM(DS:RBX &~FFFH).PENDING ≠ 0) or  
 (EPCM(DS:RBX &~FFFH).MODIFIED ≠ 0) or (EPCM(DS:RBX &~FFFH).BLOCKED ≠ 0) or  
 (EPCM(DS:RBX &~FFFH).PT ≠ PT\_REG) or (EPCM(DS:RBX &~FFFH).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RBX &~FFFH).ENCLAVEADDRESS ≠ (DS:RBX & FFFH)) )  
 THEN #PF(DS:RBX); FI;

(\* Copy 64 bytes of contents \*)  
 SCRATCH\_SECINFO ← DS:RBX;

(\* Check for mis-configured SECINFO flags\*)  
 IF (SCRATCH\_SECINFO reserved fields are not zero ) )  
 THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
 THEN #GP(0); FI;



```
IF (DS:RCX is not within CR_ELRANGE)
    THEN #GP(0); FI;
```

```
IF (DS:RCX does not resolve within an EPC)
    THEN #PF(DS:RCX); FI;
```

(\* Check that the combination of requested PT, PENDING and MODIFIED is legal \*)

```
IF (NOT (((SCRATCH_SECINFO.FLAGS.PT is PT_REG) and
    ((SCRATCH_SECINFO.FLAGS.PR is 1) or
    (SCRATCH_SECINFO.FLAGS.PENDING is 1)) and
    (SCRATCH_SECINFO.FLAGS.MODIFIED is 0)) or
    ((SCRATCH_SECINFO.FLAGS.PT is PT_TCS or PT_TRIM) and
    (SCRATCH_SECINFO.FLAGS.PR is 0) and
    (SCRATCH_SECINFO.FLAGS.PENDING is 0) and
    (SCRATCH_SECINFO.FLAGS.MODIFIED is 1) )))
    THEN #GP(0); FI
```

(\* Check security attributes of the destination EPC page \*)

```
IF ( (EPCM(DS:RCX).VALID is 0) or (EPCM(DS:RCX).BLOCKED is not 0) or
    ((EPCM(DS:RCX).PT is not PT_REG) and (EPCM(DS:RCX).PT is not PT_TCS) and (EPCM(DS:RCX).PT is not PT_TRIM)) or
    (EPCM(DS:RCX).ENCLAVESECS ≠ CR_ACTIVE_SECS))
    THEN #PF(DS:RCX); FI;
```

(\* Check the destination EPC page for concurrency \*)

```
IF ( EPC page in use )
    THEN #GP(0); FI;
```

(\* Re-Check security attributes of the destination EPC page \*)

```
IF ( (EPCM(DS:RCX).VALID is 0) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR_ACTIVE_SECS) )
    THEN #PF(DS:RCX); FI;
```

(\* Verify that accept request matches current EPC page settings \*)

```
IF ( (EPCM(DS:RCX).ENCLAVEADDRESS ≠ DS:RCX) or (EPCM(DS:RCX).PENDING ≠ SCRATCH_SECINFO.FLAGS.PENDING) or
    (EPCM(DS:RCX).MODIFIED ≠ SCRATCH_SECINFO.FLAGS.MODIFIED) or (EPCM(DS:RCX).R ≠ SCRATCH_SECINFO.FLAGS.R) or
    (EPCM(DS:RCX).W ≠ SCRATCH_SECINFO.FLAGS.W) or (EPCM(DS:RCX).X ≠ SCRATCH_SECINFO.FLAGS.X) or
    (EPCM(DS:RCX).PT ≠ SCRATCH_SECINFO.FLAGS.PT) )
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_PAGE_ATTRIBUTES_MISMATCH;
        GOTO DONE;
```

```
FI;
```

(\* Check that all required threads have left enclave \*)

```
IF (Tracking not correct)
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_NOT_TRACKED;
        GOTO DONE;
```

```
FI;
```

(\* Get pointer to the SECS to which the EPC page belongs \*)

```
TMP_SECS = << Obtain physical address of SECS through EPCM(DS:RCX)>>
```

(\* For TCS pages, perform additional checks \*)

```
IF (SCRATCH_SECINFO.FLAGS.PT = PT_TCS)
    THEN
```

## SGX INSTRUCTION REFERENCES

```
IF (DS:RCX.RESERVED ≠ 0) #GP(0); FI;  
FI;
```

(\* Check that TCS.FLAGS.DBGOPTIN, TCS stack, and TCS status are correctly initialized \*)

```
IF ( ((DS:RCX).FLAGS.DBGOPTIN is not 0) or ((DS:RCX).CSSA ≥ (DS:RCX).NSSA) or ((DS:RCX).AEP is not 0) or ((DS:RCX).STATE is not 0) )  
THEN #GP(0); FI;
```

(\* Check consistency of FS & GS Limit \*)

```
IF ( (TMP_SECS.ATTRIBUTES.MODE64BIT is 0) and ((DS:RCX.FSLIMIT & FFFH ≠ FFFH) or (DS:RCX.GSLIMIT & FFFH ≠ FFFH)) )  
THEN #GP(0); FI;
```

(\* Clear PENDING/MODIFIED flags to mark accept operation complete \*)

```
EPCM(DS:RCX).PENDING ← 0;  
EPCM(DS:RCX).MODIFIED ← 0;  
EPCM(DS:RCX).PR ← 0;
```

(\* Clear EAX and ZF to indicate successful completion \*)

```
RFLAGS.ZF ← 0;  
RAX ← 0;
```

DONE:

```
RFLAGS.CF,PF,AF,OF,SF ← 0;
```

### Flags Affected

Sets ZF if page cannot be accepted, otherwise cleared. Clears CF, PF, AF, OF, SF

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the DS segment limit.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked.       |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.<br>If EPC page has incorrect page type or security attributes. |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand is non-canonical form.<br>If a memory operand is not properly aligned.<br>If a memory operand is locked.                                   |
| #PF(error code) | If a page fault occurs in accessing memory operands.<br>If a memory operand is not an EPC page.<br>If EPC page has incorrect page type or security attributes. |

## EACCEPTCOPY—Initialize a Pending Page

| Opcode/<br>Instruction          | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|---------------------------------|-------|------------------------------|--------------------------|---|
| EAX = 07H<br>ENCLU[EACCEPTCOPY] | IR    | V/V                          | SGX2                     | This leaf function initializes a dynamically allocated EPC page from another page in the EPC. |

### Instruction Operand Encoding

| Op/En | EAX              |                         | RBX                       | RCX                                      | RDX                                 |
|-------|------------------|-------------------------|---------------------------|--|-------------------------------------|
| IR    | EACCEPTCOPY (In) | Return Error Code (Out) | Address of a SECINFO (In) | Address of the destination EPC page (In) | Address of the source EPC page (In) |

### Description

This leaf function copies the contents of an existing EPC page into an uninitialized EPC page (created by EAUG). After initialization, the instruction may also modify the access rights associated with the destination EPC page. This instruction leaf can only be executed when inside the enclave.

RBX contains the effective address of a SECINFO structure while RCX and RDX each contain the effective address of an EPC page. The table below provides additional information on the memory parameter of the EACCEPTCOPY leaf function.

### EACCEPTCOPY Memory Parameter Semantics

| SECINFO                              | EPCPAGE (Destination)                  | EPCPAGE (Source)                 |
|--------------------------------------|--|----------------------------------|
| Read access permitted by Non Enclave | Read/Write access permitted by Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

### EACCEPTCOPY Faulting Conditions

|   |  |
|---|--|
| The operands are not properly aligned.    | If security attributes of the SECINFO page make the page inaccessible.               |
| The EPC page is locked by another thread. | If security attributes of the source EPC page make the page inaccessible.            |
| The EPC page is not valid.                | RBX does not contain an effective address in an EPC page in the running enclave.     |
| SECINFO contains an invalid request.      | RCX/RDX does not contain an effective address of an EPC page in the running enclave. |

The error codes are:

**Table 40-57. EACCEPTCOPY Return Value in RAX**

| Error Code (see Table 40-4)  | Description   |
|------------------------------|---|
| No Error                     | EACCEPTCOPY successful.   |
| SGX_PAGE_ATTRIBUTES_MISMATCH | The attributes of the target EPC page do not match the expected values. |

Concurrency Restrictions

**Table 40-58. Base Concurrency Restrictions of EACCEPTCOPY**

| Leaf        | Parameter        | Base Concurrency Restrictions |             |                                    |
|-------------|------------------|-------------------------------|-------------|------------------------------------|
|             |                  | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EACCEPTCOPY | Target [DS:RCX]  | Concurrent                    |             |                                    |
|             | Source [DS:RDX]  | Concurrent                    |             |                                    |
|             | SECINFO [DS:RBX] | Concurrent                    |             |                                    |

**Table 40-59. Additional Concurrency Restrictions of EACCEPTCOPY**

| Leaf        | Parameter        | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------------|------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|             |                  | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|             |                  | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EACCEPTCOPY | Target [DS:RCX]  | Exclusive                                       | #GP         | Concurrent               |             | Concurrent          |             |
|             | Source [DS:RDX]  | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|             | SECINFO [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

**Temp Variables in EACCEPTCOPY Operational Flow**

| Name            | Type    | Size (bits) | Description   |
|-----------------|---------|-------------|---|
| SCRATCH_SECINFO | SECINFO | 512         | Scratch storage for holding the contents of DS:RBX. |

IF (DS:RBX is not 64Byte Aligned)  
 THEN #GP(0); FI;

IF ( (DS:RCX is not 4KByte Aligned) or (DS:RDX is not 4KByte Aligned) )  
 THEN #GP(0); FI;

IF ((DS:RBX is not within CR\_ELRANGE) or (DS:RCX is not within CR\_ELRANGE) or (DS:RDX is not within CR\_ELRANGE))  
 THEN #GP(0); FI;

IF (DS:RBX does not resolve within an EPC)  
 THEN #PF(DS:RBX); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

IF (DS:RDX does not resolve within an EPC)  
 THEN #PF(DS:RDX); FI;

IF ( (EPCM(DS:RBX &~FFFH).VALID = 0) or (EPCM(DS:RBX &~FFFH).R = 0) or (EPCM(DS:RBX &~FFFH).PENDING ≠ 0) or  
 (EPCM(DS:RBX &~FFFH).MODIFIED ≠ 0) or (EPCM(DS:RBX &~FFFH).BLOCKED ≠ 0) or (EPCM(DS:RBX &~FFFH).PT ≠ PT\_REG) or  
 (EPCM(DS:RBX &~FFFH).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RBX &~FFFH).ENCLAVEADDRESS ≠ DS:RBX) )  
 THEN #PF(DS:RBX); FI;

(\* Copy 64 bytes of contents \*)  
 SCRATCH\_SECINFO ← DS:RBX;

(\* Check for mis-configured SECINFO flags\*)  
 IF ( (SCRATCH\_SECINFO reserved fields are not zero ) or ((SCRATCH\_SECINFO.FLAGS.R=0) AND(SCRATCH\_SECINFO.FLAGS.W≠0 ) or  
 (SCRATCH\_SECINFO.FLAGS.PT is not PT\_REG) )  
 THEN #GP(0); FI;

(\* Check security attributes of the source EPC page \*)  
 IF ( (EPCM(DS:RDX).VALID = 0) or (EPCM(DS:RCX).R = 0) or (EPCM(DS:RDX).PENDING ≠ 0) or (EPCM(DS:RDX).MODIFIED ≠ 0) or  
 (EPCM(DS:RDX).BLOCKED ≠ 0) or (EPCM(DS:RDX).PT ≠ PT\_REG) or (EPCM(DS:RDX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RDX).ENCLAVEADDRESS ≠ DS:RDX))  
 THEN #PF(DS:RDX); FI;

(\* Check security attributes of the destination EPC page \*)  
 IF ( (EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PENDING ≠ 1) or (EPCM(DS:RCX).MODIFIED ≠ 0) or  
 (EPCM(DS:RDX).BLOCKED ≠ 0) or (EPCM(DS:RCX).PT ≠ PT\_REG) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) )  
 THEN  
 RFLAGS.ZF ← 1;  
 RAX ← SGX\_PAGE\_ATTRIBUTES\_MISMATCH;  
 GOTO DONE;  
 FI;

(\* Check the destination EPC page for concurrency \*)  
 IF (destination EPC page in use )  
 THEN #GP(0); FI;

(\* Re-Check security attributes of the destination EPC page \*)  
 IF ( (EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PENDING ≠ 1) or (EPCM(DS:RCX).MODIFIED ≠ 0) or  
 (EPCM(DS:RCX).R ≠ 1) or (EPCM(DS:RCX).W ≠ 1) or (EPCM(DS:RCX).X ≠ 0) or  
 (EPCM(DS:RCX).PT ≠ SCRATCH\_SECINFO.FLAGS.PT) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RCX).ENCLAVEADDRESS ≠ DS:RCX))  
 THEN  
 RFLAGS.ZF ← 1;  
 RAX ← SGX\_PAGE\_ATTRIBUTES\_MISMATCH;  
 GOTO DONE;  
 FI;

(\* Copy 4Kbytes form the source to destination EPC page\*)  
 DS:RCX[32767:0] ← DS:RDX[32767:0];

(\* Update EPCM permissions \*)  
 EPCM(DS:RCX).R ← SCRATCH\_SECINFO.FLAGS.R;  
 EPCM(DS:RCX).W ← SCRATCH\_SECINFO.FLAGS.W;  
 EPCM(DS:RCX).X ← SCRATCH\_SECINFO.FLAGS.X;  
 EPCM(DS:RCX).PENDING ← 0;

RFLAGS.ZF ← 0;  
 RAX ← 0;

DONE:  
 RFLAGS.CF,PF,AF,OF,SF ← 0;

**Flags Affected**

Sets ZF if page is not modifiable, otherwise cleared. Clears CF, PF, AF, OF, SF

**Protected Mode Exceptions**

- #GP(0) If a memory operand effective address is outside the DS segment limit.  
If a memory operand is not properly aligned.  
If a memory operand is locked.
- #PF(error code) If a page fault occurs in accessing memory operands.  
If a memory operand is not an EPC page.  
If EPC page has incorrect page type or security attributes.

**64-Bit Mode Exceptions**

- #GP(0) If a memory operand is non-canonical form.  
If a memory operand is not properly aligned.  
If a memory operand is locked.
- #PF(error code) If a page fault occurs in accessing memory operands.  
If a memory operand is not an EPC page.  
If EPC page has incorrect page type or security attributes.

## EENTER—Enters an Enclave

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                     |
|----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 02H<br>ENCLU[EENTER] | IR    | V/V                          | SGX1                     | This leaf function is used to enter an enclave. |

### Instruction Operand Encoding

| Op/En | EAX         |                              | RBX                   | RCX                 |   |
|-------|-------------|------------------------------|-----------------------|---------------------|---|
| IR    | EENTER (In) | Content of RBX.CSSA<br>(Out) | Address of a TCS (In) | Address of AEP (In) | Address of IP following<br>EENTER (Out) |

### Description

The ENCLU[EENTER] instruction transfers execution to an enclave. At the end of the instruction, the logical processor is executing in enclave mode at the RIP computed as EnclaveBase + TCS.OENTRY. If the target address is not within the CS segment (32-bit) or is not canonical (64-bit), a #GP(0) results.

### EENTER Memory Parameter Semantics

|                |
|----------------|
| TCS            |
| Enclave access |

EENTER is a serializing instruction. The instruction faults if any of the following occurs:

|  |   |
|--|---|
| Address in RBX is not properly aligned.                    | Any TCS.FLAGS's must-be-zero bit is not zero.   |
| TCS pointed to by RBX is not valid or available or locked. | Current 32/64 mode does not match the enclave mode in SECS.ATTRIBUTES.MODE64.         |
| The SECS is in use.  | Either of TCS-specified FS and GS segment is not a subsets of the current DS segment. |
| Any one of DS, ES, CS, SS is not zero.                     | If XSAVE available, CR4.OSXSAVE = 0, but SECS.ATTRIBUTES.XFRM ≠ 3.                    |
| CR4.OSFXSR ≠ 1.  | If CR4.OSXSAVE = 1, SECS.ATTRIBUTES.XFRM is not a subset of XCR0.                     |

The following operations are performed by EENTER:

- RSP and RBP are saved in the current SSA frame on EENTER and are automatically restored on EEXIT or interrupt.
- The AEP contained in RCX is stored into the TCS for use by AEXs. FS and GS (including hidden portions) are saved and new values are constructed using TCS.OFSBASE/GSBASE (32 and 64-bit mode) and TCS.OFSLIMIT/GSLIMIT (32-bit mode only). The resulting segments must be a subset of the DS segment.
- If CR4.OSXSAVE == 1, XCR0 is saved and replaced by SECS.ATTRIBUTES.XFRM. The effect of RFLAGS.TF depends on whether the enclave entry is opt-in or opt-out (see Section 42.1.2):
  - On opt-out entry, TF is saved and cleared (it is restored on EEXIT or AEX). Any attempt to set TF via a POPF instruction while inside the enclave clears TF (see Section 42.2.5).
  - On opt-in entry, a single-step debug exception is pending on the instruction boundary immediately after EENTER (see Section 42.2.2).
- All code breakpoints that do not overlap with ELRANGE are also suppressed. If the entry is an opt-out entry, all code and data breakpoints that overlap with the ELRANGE are suppressed.
- On opt-out entry, a number of performance monitoring counters and behaviors are modified or suppressed (see Section 42.2.3):

- All performance monitoring activity on the current thread is suppressed except for incrementing and firing of FIXED\_CTR1 and FIXED\_CTR2.
- PEBS is suppressed.
- AnyThread counting on other threads is demoted to MyThread mode and IA32\_PERF\_GLOBAL\_STATUS[60] on that thread is set
- If the opt-out entry on a hardware thread results in suppression of any performance monitoring, then the processor sets IA32\_PERF\_GLOBAL\_STATUS[60] and IA32\_PERF\_GLOBAL\_STATUS[63].

Concurrency Restrictions

Table 40-60. Base Concurrency Restrictions of EENTER

| Leaf   | Parameter    | Base Concurrency Restrictions |             |                                    |
|--------|--------------|-------------------------------|-------------|------------------------------------|
|        |              | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EENTER | TCS [DS:RBX] | Shared                        | #GP         |                                    |

Table 40-61. Additional Concurrency Restrictions of EENTER

| Leaf   | Parameter    | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------|--------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|        |              | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |              | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EENTER | TCS [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

Temp Variables in EENTER Operational Flow

| Name         | Type              | Size (Bits) | Description  |
|--------------|-------------------|-------------|--|
| TMP_FSBASE   | Effective Address | 32/64       | Proposed base address for FS segment.                            |
| TMP_GSBASE   | Effective Address | 32/64       | Proposed base address for GS segment.                            |
| TMP_FSLIMIT  | Effective Address | 32/64       | Highest legal address in proposed FS segment.                    |
| TMP_GSLIMIT  | Effective Address | 32/64       | Highest legal address in proposed GS segment.                    |
| TMP_XSIZE    | integer           | 64          | Size of XSAVE area based on SECS.ATTRIBUTES.XFRM.                |
| TMP_SSA_PAGE | Effective Address | 32/64       | Pointer used to iterate over the SSA pages in the current frame. |
| TMP_GPR      | Effective Address | 32/64       | Address of the GPR area within the current SSA frame.            |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

(\* Make sure DS is usable, expand up \*)

IF (TMP\_MODE64 = 0 and (DS not usable or ( ( DS[S] = 1) and (DS[bit 11] = 0) and DS[bit 10] = 1) ) )  
 THEN #GP(0); FI;

(\* Check that CS, SS, DS, ES.base is 0 \*)

IF (TMP\_MODE64 = 0)  
 THEN  
 IF(CS.base ≠ 0 or DS.base ≠ 0) #GP(0); FI;  
 IF(ES usable and ES.base ≠ 0) #GP(0); FI;  
 IF(SS usable and SS.base ≠ 0) #GP(0); FI;  
 IF(SS usable and SS.B = 0) #GP(0); FI;



```

FI;

IF (DS:RBX is not 4KByte Aligned)
  THEN #GP(0); FI;

IF (DS:RBX does not resolve within an EPC)
  THEN #PF(DS:RBX); FI;

(* Check AEP is canonical*)
IF (TMP_MODE64 = 1 and (CS:RCX is not canonical) )
  THEN #GP(0); FI;

(* Check concurrency of TCS operation*)
IF (Other Intel SGX instructions is operating on TCS)
  THEN #GP(0); FI;

(* TCS verification *)
IF (EPCM(DS:RBX).VALID = 0)
  THEN #PF(DS:RBX); FI;

IF (EPCM(DS:RBX).BLOCKED = 1)
  THEN #PF(DS:RBX); FI;

IF ( ( EPCM(DS:RBX).ENCLAVEADDRESS ≠ DS:RBX) or (EPCM(DS:RBX).PT ≠ PT_TCS) )
  THEN #PF(DS:RBX); FI;

IF ((EPCM(DS:RBX).PENDING = 1) or (EPCM(DS:RBX).MODIFIED = 1))
  THEN #PF(DS:RBX); FI;

IF ( (DS:RBX).OSSA is not 4KByte Aligned)
  THEN #GP(0); FI;

(* Check proposed FS and GS *)
IF ( ( (DS:RBX).OFSBASE is not 4KByte Aligned) or ( (DS:RBX).OGSBASE is not 4KByte Aligned) )
  THEN #GP(0); FI;

(* Get the SECS for the enclave in which the TCS resides *)
TMP_SECS ← Address of SECS for TCS;

(* Check proposed FS/GS segments fall within DS *)
IF (TMP_MODE64 = 0)
  THEN
    TMP_FSBASE ← (DS:RBX).OFSBASE + TMP_SECS.BASEADDR;
    TMP_FSLIMIT ← (DS:RBX).OFSBASE + TMP_SECS.BASEADDR + (DS:RBX).FSLIMIT;
    TMP_GSBASE ← (DS:RBX).OGSBASE + TMP_SECS.BASEADDR;
    TMP_GSLIMIT ← (DS:RBX).OGSBASE + TMP_SECS.BASEADDR + (DS:RBX).GSLIMIT;
    (* if FS wrap-around, make sure DS has no holes*)
    IF (TMP_FSLIMIT < TMP_FSBASE)
      THEN
        IF (DS.limit < 4GB) THEN #GP(0); FI;
        ELSE
          IF (TMP_FSLIMIT > DS.limit) THEN #GP(0); FI;
      FI;
    (* if GS wrap-around, make sure DS has no holes*)

```

## SGX INSTRUCTION REFERENCES

```
    IF (TMP_GSLIMIT < TMP_GSBASE)
        THEN
            IF (DS.limit < 4GB) THEN #GP(0); FI;
        ELSE
            IF (TMP_GSLIMIT > DS.limit) THEN #GP(0); FI;
    FI;
ELSE
    TMP_FSBASE ← (DS:RBX).OFSBASE + TMP_SECS.BASEADDR;
    TMP_GSBASE ← (DS:RBX).OGSBASE + TMP_SECS.BASEADDR;
    IF ( (TMP_FSBASE is not canonical) or (TMP_GSBASE is not canonical))
        THEN #GP(0); FI;
FI;

(* Ensure that the FLAGS field in the TCS does not have any reserved bits set *)
IF ( ( (DS:RBX).FLAGS & FFFFFFFF7FFF) ≠ 0 )
    THEN #GP(0); FI;

(* SECS must exist and enclave must have previously been EINITted *)
IF (the enclave is not already initialized)
    THEN #GP(0); FI;

(* make sure the logical processor's operating mode matches the enclave *)
IF ( (TMP_MODE64 ≠ TMP_SECS.ATTRIBUTES.MODE64BIT) )
    THEN #GP(0); FI;

IF (CR4.OSFXSR = 0)
    THEN #GP(0); FI;

(* Check for legal values of SECS.ATTRIBUTES.XFRM *)
IF (CR4.OSXSAVE = 0)
    THEN
        IF (TMP_SECS.ATTRIBUTES.XFRM ≠ 03H) THEN #GP(0); FI;
    ELSE
        IF ( (TMP_SECS.ATTRIBUTES.XFRM & XCRO) ≠ TMP_SECS.ATTRIBUTES.XFRM) THEN #GP(0); FI;
FI;

(* Make sure the SSA contains at least one more frame *)
IF ( (DS:RBX).CSSA ≥ (DS:RBX).NSSA)
    THEN #GP(0); FI;

(* Compute linear address of SSA frame *)
TMP_SSA ← (DS:RBX).OSSA + TMP_SECS.BASEADDR + 4096 * TMP_SECS.SSAFRAMESIZE * (DS:RBX).CSSA;
TMP_XSIZE ← compute_XSAVE_frame_size(TMP_SECS.ATTRIBUTES.XFRM);

FOR EACH TMP_SSA_PAGE = TMP_SSA to TMP_SSA + TMP_XSIZE
    (* Check page is read/write accessible *)
    Check that DS:TMP_SSA_PAGE is read/write accessible;
    If a fault occurs, release locks, abort and deliver that fault;

    IF (DS:TMP_SSA_PAGE does not resolve to EPC page)
        THEN #PF(DS:TMP_SSA_PAGE); FI;
    IF (EPCM(DS:TMP_SSA_PAGE).VALID = 0)
        THEN #PF(DS:TMP_SSA_PAGE); FI;
    IF (EPCM(DS:TMP_SSA_PAGE).BLOCKED = 1)
```

```

    THEN #PF(DS:TMP_SSA_PAGE); FI;
  IF ((EPCM(DS:TMP_SSA_PAGE).PENDING = 1) or (EPCM(DS:TMP_SSA_PAGE).MODIFIED = 1))
    THEN #PF(DS:TMP_SSA_PAGE); FI;
  IF ( ( EPCM(DS:TMP_SSA_PAGE).ENCLAVEADDRESS ≠ DS:TMP_SSA_PAGE) or (EPCM(DS:TMP_SSA_PAGE).PT ≠ PT_REG) or
    (EPCM(DS:TMP_SSA_PAGE).ENCLAVESECS ≠ EPCM(DS:RBX).ENCLAVESECS) or
    (EPCM(DS:TMP_SSA_PAGE).R = 0) or (EPCM(DS:TMP_SSA_PAGE).W = 0) )
    THEN #PF(DS:TMP_SSA_PAGE); FI;
  CR_XSAVE_PAGE_n ← Physical_Address(DS:TMP_SSA_PAGE);
ENDFOR

```

(\* Compute address of GPR area\*)

```
TMP_GPR ← TMP_SSA + 4096 * DS:TMP_SECS.SSAFRAMESIZE -- sizeof(GPRSGX_AREA);
```

If a fault occurs; release locks, abort and deliver that fault;

```

IF (DS:TMP_GPR does not resolve to EPC page)
  THEN #PF(DS:TMP_GPR); FI;
IF (EPCM(DS:TMP_GPR).VALID = 0)
  THEN #PF(DS:TMP_GPR); FI;
IF (EPCM(DS:TMP_GPR).BLOCKED = 1)
  THEN #PF(DS:TMP_GPR); FI;
IF ((EPCM(DS:TMP_GPR).PENDING = 1) or (EPCM(DS:TMP_GPR).MODIFIED = 1))
  THEN #PF(DS:TMP_GPR); FI;
IF ( ( EPCM(DS:TMP_GPR).ENCLAVEADDRESS ≠ DS:TMP_GPR) or (EPCM(DS:TMP_GPR).PT ≠ PT_REG) or
  (EPCM(DS:TMP_GPR).ENCLAVESECS ≠ EPCM(DS:RBX).ENCLAVESECS) or
  (EPCM(DS:TMP_GPR).R = 0) or (EPCM(DS:TMP_GPR).W = 0) )
  THEN #PF(DS:TMP_GPR); FI;

```

```

IF (TMP_MODE64 = 0)
  THEN
    IF (TMP_GPR + (GPR_SIZE - 1) is not in DS segment) THEN #GP(0); FI;
FI;

```

```
CR_GPR_PA ← Physical_Address (DS: TMP_GPR);
```

(\* Validate TCS.OENTRY \*)

```
TMP_TARGET ← (DS:RBX).OENTRY + TMP_SECS.BASEADDR;
```

```

IF (TMP_MODE64 = 1)
  THEN
    IF (TMP_TARGET is not canonical) THEN #GP(0); FI;
  ELSE
    IF (TMP_TARGET > CS limit) THEN #GP(0); FI;
FI;

```

(\* Ensure the enclave is not already active and this thread is the only one using the TCS\*)

```

IF (DS:RBX.STATE = ACTIVE)
  THEN #GP(0); FI;

```

```

CR_ENCLAVE_MODE ← 1;
CR_ACTIVE_SECS ← TMP_SECS;
CR_ELRRANGE ← (TMPSECS.BASEADDR, TMP_SECS.SIZE);

```

(\* Save state for possible AEXs \*)

```
CR_TCS_PA ← Physical_Address (DS:RBX);
```

```
CR_TCS_LA ← RBX;
```

## SGX INSTRUCTION REFERENCES

CR\_TCS\_LA.AEP ← RCX;

(\* Save the hidden portions of FS and GS \*)

CR\_SAVE\_FS\_selector ← FS.selector;

CR\_SAVE\_FS\_base ← FS.base;

CR\_SAVE\_FS\_limit ← FS.limit;

CR\_SAVE\_FS\_access\_rights ← FS.access\_rights;

CR\_SAVE\_GS\_selector ← GS.selector;

CR\_SAVE\_GS\_base ← GS.base;

CR\_SAVE\_GS\_limit ← GS.limit;

CR\_SAVE\_GS\_access\_rights ← GS.access\_rights;

(\* If XSAVE is enabled, save XCRO and replace it with SECS.ATTRIBUTES.XFRM\*)

IF (CR4.OSXSAVE = 1)

    CR\_SAVE\_XCRO ← XCRO;

    XCRO ← TMP\_SECS.ATTRIBUTES.XFRM;

FI;

RCX ← RIP;

RIP ← TMP\_TARGET;

RAX ← (DS:RBX).CSSA;

(\* Save the outside RSP and RBP so they can be restored on interrupt or EEXIT \*)

DS:TMP\_SSA.U\_RSP ← RSP;

DS:TMP\_SSA.U\_RBP ← RBP;

(\* Do the FS/GS swap \*)

FS.base ← TMP\_FSBASE;

FS.limit ← DS:RBX.FSLIMIT;

FS.type ← 0001b;

FS.W ← DS.W;

FS.S ← 1;

FS.DPL ← DS.DPL;

FS.G ← 1;

FS.B ← 1;

FS.P ← 1;

FS.AVL ← DS.AVL;

FS.L ← DS.L;

FS.unusable ← 0;

FS.selector ← 0BH;

GS.base ← TMP\_GSBASE;

GS.limit ← DS:RBX.GSLIMIT;

GS.type ← 0001b;

GS.W ← DS.W;

GS.S ← 1;

GS.DPL ← DS.DPL;

GS.G ← 1;

GS.B ← 1;

GS.P ← 1;

GS.AVL ← DS.AVL;

GS.L ← DS.L;

GS.unusable ← 0;

GS.selector ← 0BH;

```

CR_DBGOPTIN ← TCS.FLAGS.DBGOPTIN;
Suppress_all_code_breakpoints_that_are_outside_ELRANGE;

IF (CR_DBGOPTIN = 0)
  THEN
    Suppress_all_code_breakpoints_that_overlap_with_ELRANGE;
    CR_SAVE_TF ← RFLAGS.TF;
    RFLAGS.TF ← 0;
    Suppress_monitor_trap_flag for the source of the execution of the enclave;
    Suppress any pending debug exceptions;
    Suppress any pending MTF VM exit;
  ELSE
    IF RFLAGS.TF = 1
      THEN pend a single-step #DB at the end of EENTER; FI;
    IF the "monitor trap flag" VM-execution control is set
      THEN pend an MTF VM exit at the end of EENTER; FI;
  FI;

Flush_linear_context;
Allow_front_end_to_begin_fetch_at_new_RIP;

```

### Flags Affected

RFLAGS.TF is cleared on opt-out entry

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <ul style="list-style-type: none"> <li>If DS:RBX is not page aligned.</li> <li>If the enclave is not initialized.</li> <li>If part or all of the FS or GS segment specified by TCS is outside the DS segment or not properly aligned.</li> <li>If the thread is not in the INACTIVE state.</li> <li>If CS, DS, ES or SS bases are not all zero.</li> <li>If executed in enclave mode.</li> <li>If any reserved field in the TCS FLAG is set.</li> <li>If the target address is not within the CS segment.</li> <li>If CR4.OSFXSR = 0.</li> <li>If CR4.OSXSAVE = 0 and SECS.ATTRIBUTES.XFRM ≠ 3.</li> <li>If CR4.OSXSAVE = 1 and SECS.ATTRIBUTES.XFRM is not a subset of XCRO.</li> </ul> |
| #PF(error code) | <ul style="list-style-type: none"> <li>If a page fault occurs in accessing memory.</li> <li>If DS:RBX does not point to a valid TCS.</li> <li>If one or more pages of the current SSA frame are not readable/writable, or do not resolve to a valid PT_REG EPC page.</li> </ul>  |

**64-Bit Mode Exceptions**

- #GP(0)
  - If DS:RBX is not page aligned.
  - If the enclave is not initialized.
  - If the thread is not in the INACTIVE state.
  - If CS, DS, ES or SS bases are not all zero.
  - If executed in enclave mode.
  - If part or all of the FS or GS segment specified by TCS is outside the DS segment or not properly aligned.
  - If the target address is not canonical.
  - If CR4.OSFXSR = 0.
  - If CR4.OSXSAVE = 0 and SECS.ATTRIBUTES.XFRM ≠ 3.
  - If CR4.OSXSAVE = 1 and SECS.ATTRIBUTES.XFRM is not a subset of XCR0.
- #PF(error code)
  - If a page fault occurs in accessing memory operands.
  - If DS:RBX does not point to a valid TCS.
  - If one or more pages of the current SSA frame are not readable/writable, or do not resolve to a valid PT\_REG EPC page.

## EEXIT—Exits an Enclave

| Opcode/<br>Instruction    | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                    |
|---------------------------|-------|------------------------------|--------------------------|--|
| EAX = 04H<br>ENCLU[EEXIT] | IR    | V/V                          | SGX1                     | This leaf function is used to exit an enclave. |

### Instruction Operand Encoding

| Op/En | EAX        | RBX                                     | RCX                             |
|-------|------------|---|---------------------------------|
| IR    | EEXIT (In) | Target address outside the enclave (In) | Address of the current AEP (In) |

### Description

The ENCLU[EEXIT] instruction exits the currently executing enclave and branches to the location specified in RBX. RCX receives the current AEP. If RBX is not within the CS (32-bit mode) or is not canonical (64-bit mode) a #GP(0) results.

### EEXIT Memory Parameter Semantics

| Target Address                      |
|-------------------------------------|
| Non-Enclave read and execute access |

If RBX specifies an address that is inside the enclave, the instruction will complete normally. The fetch of the next instruction will occur in non-enclave mode, but will attempt to fetch from inside the enclave. This fetch returns a fixed data pattern.

If secrets are contained in any registers, it is responsibility of enclave software to clear those registers.

If XCR0 was modified on enclave entry, it is restored to the value it had at the time of the most recent EENTER or ERESUME.

If the enclave is opt-out, RFLAGS.TF is loaded from the value previously saved on EENTER.

Code and data breakpoints are unsuppressed.

Performance monitoring counters are unsuppressed.

### Concurrency Restrictions

**Table 40-62. Base Concurrency Restrictions of EEXIT**

| Leaf  | Parameter | Base Concurrency Restrictions |             |                                    |
|-------|-----------|-------------------------------|-------------|------------------------------------|
|       |           | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EEXIT |           | Concurrent                    |             |                                    |

**Table 40-63. Additional Concurrency Restrictions of EEXIT**

| Leaf  | Parameter | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------|-----------|---|-------------|--------------------------|-------------|---------------------|-------------|
|       |           | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|       |           | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EEXIT |           | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

Temp Variables in EEXIT Operational Flow

| Name    | Type              | Size (Bits) | Description                                   |
|---------|-------------------|-------------|---|
| TMP_RIP | Effective Address | 32/64       | Saved copy of CRIP for use when creating LBR. |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

```
IF (TMP_MODE64 = 1)
    THEN
        IF (RBX is not canonical) THEN #GP(0); FI;
    ELSE
        IF (RBX > CS limit) THEN #GP(0); FI;
FI;
```

TMP\_RIP ← CRIP;  
RIP ← RBX;

(\* Return current AEP in RCX \*)  
RCX ← CR\_TCS\_PA.AEP;

(\* Do the FS/GS swap \*)  
FS.selector ← CR\_SAVE\_FS.selector;  
FS.base ← CR\_SAVE\_FS.base;  
FS.limit ← CR\_SAVE\_FS.limit;  
FS.access\_rights ← CR\_SAVE\_FS.access\_rights;  
GS.selector ← CR\_SAVE\_GS.selector;  
GS.base ← CR\_SAVE\_GS.base;  
GS.limit ← CR\_SAVE\_GS.limit;  
GS.access\_rights ← CR\_SAVE\_GS.access\_rights;

(\* Restore XCRO if needed \*)  
IF (CR4.OSXSAVE = 1)  
 XCRO ← CR\_SAVE\_\_XCRO;  
FI;

Unsuppress\_all\_code\_breakpoints\_that\_are\_outside\_ELRange;

```
IF (CR_DBGOPTIN = 0)
    THEN
        UnSuppress_all_code_breakpoints_that_overlap_with_ELRange;
        Restore suppressed breakpoint matches;
        RFLAGS.TF ← CR_SAVE_TF;
        UnSuppress_monitor_trap_flag;
        UnSuppress_LBR_Generation;
        UnSuppress_performance_monitoring_activity;
        Restore performance monitoring counter AnyThread demotion to MyThread in enclave back to AnyThread
FI;
```

```
IF RFLAGS.TF = 1
    THEN Pend Single-Step #DB at the end of EEXIT;
FI;
```



IF the “monitor trap flag” VM-execution control is set  
 THEN pend a MTF VM exit at the end of EEXIT;  
 FI;

CR\_ENCLAVE\_MODE ← 0;  
 CR\_TCS\_PA.STATE ← INACTIVE;

(\* Assure consistent translations \*)  
 Flush\_linear\_context;

### Flags Affected

RFLAGS.TF is restored from the value previously saved in EENTER or ERESUME.

### Protected Mode Exceptions

#GP(0)                If executed outside an enclave.  
                           If RBX is outside the CS segment.  
 #PF(error code)    If a page fault occurs in accessing memory.

### 64-Bit Mode Exceptions

#GP(0)                If executed outside an enclave.  
                           If RBX is not canonical.  
 #PF(error code)    If a page fault occurs in accessing memory operands.

## EGETKEY—Retrieves a Cryptographic Key

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description                                       |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 01H<br>ENCLU[EGETKEY] | IR    | V/V                          | SGX1                     | This leaf function retrieves a cryptographic key. |

### Instruction Operand Encoding

| Op/En | EAX          | RBX                          | RCX                            |
|-------|--------------|------------------------------|--------------------------------|
| IR    | EGETKEY (In) | Address to a KEYREQUEST (In) | Address of the OUTPUTDATA (In) |

### Description

The ENCLU[EGETKEY] instruction returns a 128-bit secret key from the processor specific key hierarchy. The register RBX contains the effective address of a KEYREQUEST structure, which the instruction interprets to determine the key being requested. The Requesting Keys section below provides a description of the keys that can be requested. The RCX register contains the effective address where the key will be returned. Both the addresses in RBX & RCX should be locations inside the enclave.

EGETKEY derives keys using a processor unique value to create a specific key based on a number of possible inputs. This instruction leaf can only be executed inside an enclave.

### EGETKEY Memory Parameter Semantics

| KEYREQUEST          | OUTPUTDATA           |
|---------------------|----------------------|
| Enclave read access | Enclave write access |

After validating the operands, the instruction determines which key is to be produced and performs the following actions:

- The instruction assembles the derivation data for the key based on the Table 40-64.
- Computes derived key using the derivation data and package specific value.
- Outputs the calculated key to the address in RCX.

The instruction fails with #GP(0) if the operands are not properly aligned. Successful completion of the instruction will clear RFLAGS.{ZF, CF, AF, OF, SF, PF}. The instruction returns an error code if the user tries to request a key based on an invalid CPUSVN or ISVSVN (when the user request is accepted, see the table below), requests a key for which it has not been granted the attribute to request, or requests a key that is not supported by the hardware. These checks may be performed in any order. Thus, an indication by error number of one cause (for example, invalid attribute) does not imply that there are not also other errors. Different processors may thus give different error numbers for the same Enclave. The correctness of software should not rely on the order resulting from the checks documented in this section. In such cases the ZF flag is set and the corresponding error bit (SGX\_INVALID\_SVN, SGX\_INVALID\_ATTRIBUTE, SGX\_INVALID\_KEYNAME) is set in RAX and the data at the address specified by RCX is unmodified.

### Requesting Keys

The KEYREQUEST structure (see Section 37.17.1) identifies the key to be provided. The Keyrequest.KeyName field identifies which type of key is requested.

### Deriving Keys

Key derivation is based on a combination of the enclave specific values (see Table 40-64) and a processor key. Depending on the key being requested a field may either be included by definition or the value may be included from the KeyRequest. A “yes” in Table 40-64 indicates the value for the field is included from its default location, identified in the source row, and a “request” indicates the values for the field is included from its corresponding KeyRequest field.

Table 40-64. Key Derivation

|                   | Key Name               | Attributes   | Owner Epoch        | CPU SVN             | ISV SVN         | ISV PROPID  | ISVEXT PROPID      | ISVFAM ILYID      | MRENCLAVE       | MRSIGNER       | CONFIG ID      | CONFIGS VN      | RAND       |
|-------------------|------------------------|--|--------------------|---------------------|-----------------|-------------|--------------------|-------------------|-----------------|----------------|----------------|-----------------|------------|
| Source            | Key Dependent Constant | Y← SECS.ATTRIBUTES and SECS.MISCSELECT;            | CR_SGX OWNER EPOCH | Y← CPUSVN Register; | R← Req.ISV SVN; | SECS. ISVID | SECS.IS VEXTPROPID | SECS.IS VFAMILYID | SECS. MRENCLAVE | SECS. MRSIGNER | SECS.CO NFIGID | SECS.CO NFIGSVN | Req. KEYID |
|                   |                        | R← AttrMask & SECS.ATTRIBUTES and SECS.MISCSELECT; |                    | R← Req.CPU SVN;     |                 |             |                    |                   |                 |                |                |                 |            |
| EINITTOKEN        | Yes                    | Request  | Yes                | Request             | Request         | Yes         | No                 | No                | No              | Yes            | No             | No              | Request    |
| Report            | Yes                    | Yes  | Yes                | Yes                 | No              | No          | No                 | No                | Yes             | No             | Yes            | Yes             | Request    |
| Seal              | Yes                    | Request  | Yes                | Request             | Request         | Request     | Request            | Request           | Request         | Request        | Request        | Request         | Request    |
| Provisioning      | Yes                    | Request  | No                 | Request             | Request         | Yes         | No                 | No                | No              | Yes            | No             | No              | Yes        |
| Provisioning Seal | Yes                    | Request  | No                 | Request             | Request         | Request     | Request            | Request           | No              | Yes            | Request        | Request         | Yes        |

Keys that permit the specification of a CPU or ISV's code's, or enclave configuration's SVNs have additional requirements. The caller may not request a key for an SVN beyond the current CPU, ISV or enclave configuration's SVN, respectively.

Several keys are access controlled. Access to the Provisioning Key and Provisioning Seal key requires the enclave's ATTRIBUTES.PROVISIONKEY be set. The EINITTOKEN Key requires ATTRIBUTES.EINITTOKEN\_KEY be set and SECS.MRSIGNER equal IA32\_SGXLEPUBKEYHASH.

Some keys are derived based on a hardcoded PKCS padding constant (352 byte string):

HARDCODED\_PKCS1\_5\_PADDING[15:0] ← 0100H;

HARDCODED\_PKCS1\_5\_PADDING[2655:16] ← SignExtend330Byte(-1); // 330 bytes of 0FFH

HARDCODED\_PKCS1\_5\_PADDING[2815:2656] ← 2004000501020403650148866009060D30313000H;

The error codes are:

Table 40-65. EGETKEY Return Value in RAX

| Error Code (see Table 40-4) | Value | Description   |
|-----------------------------|-------|---|
| No Error                    | 0     | EGETKEY successful.   |
| SGX_INVALID_ATTRIBUTE       |       | The KEYREQUEST contains a KEYNAME for which the enclave is not authorized.                        |
| SGX_INVALID_CPUSVN          |       | If KEYREQUEST.CPUSVN is an unsupported platforms CPUSVN value.                                    |
| SGX_INVALID_ISVSVN          |       | If KEYREQUEST software SVN (ISVSVN or CONFIGSVN) is greater than the enclave's corresponding SVN. |
| SGX_INVALID_KEYNAME         |       | If KEYREQUEST.KEYNAME is an unsupported value.  |

## Concurrency Restrictions

Table 40-66. Base Concurrency Restrictions of EGETKEY

| Leaf    | Parameter           | Base Concurrency Restrictions |             |                                    |
|---------|---------------------|-------------------------------|-------------|------------------------------------|
|         |                     | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EGETKEY | KEYREQUEST [DS:RBX] | Concurrent                    |             |                                    |
|         | OUTPUTDATA [DS:RCX] | Concurrent                    |             |                                    |

**Table 40-67. Additional Concurrency Restrictions of EGETKEY**

| Leaf    | Parameter           | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|---------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                     | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                     | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EGETKEY | KEYREQUEST [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|         | OUTPUTDATA [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EGETKEY Operational Flow**

| Name                | Type | Size (Bits) | Description  |
|---------------------|------|-------------|--|
| TMP_CURRENTSECS     |      |             | Address of the SECS for the currently executing enclave.   |
| TMP_KEYDEPENDENCIES |      |             | Temp space for key derivation.                             |
| TMP_ATTRIBUTES      |      | 128         | Temp Space for the calculation of the sealable Attributes. |
| TMP_ISVEXTPRODID    |      | 16 bytes    | Temp Space for ISVEXTPRODID.                               |
| TMP_ISVPRODID       |      | 2 bytes     | Temp Space for ISVPRODID.                                  |
| TMP_ISVFAMILYID     |      | 16 bytes    | Temp Space for ISVFAMILYID.                                |
| TMP_CONFIGID        |      | 64 bytes    | Temp Space for CONFIGID.                                   |
| TMP_CONFIGSVN       |      | 2 bytes     | Temp Space for CONFIGSVN.                                  |
| TMP_OUTPUTKEY       |      | 128         | Temp Space for the calculation of the key.                 |

(\* Make sure KEYREQUEST is properly aligned and inside the current enclave \*)

IF ( (DS:RBX is not 512Byte aligned) or (DS:RBX is within CR\_ELRANGE) )  
 THEN #GP(0); FI;

(\* Make sure DS:RBX is an EPC address and the EPC page is valid \*)

IF ( (DS:RBX does not resolve to an EPC address) or (EPCM(DS:RBX).VALID = 0) )  
 THEN #PF(DS:RBX); FI;

IF (EPCM(DS:RBX).BLOCKED = 1)

THEN #PF(DS:RBX); FI;

(\* Check page parameters for correctness \*)

IF ( (EPCM(DS:RBX).PT ≠ PT\_REG) or (EPCM(DS:RBX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or (EPCM(DS:RBX).PENDING = 1) or  
 (EPCM(DS:RBX).MODIFIED = 1) or (EPCM(DS:RBX).ENCLAVEADDRESS ≠ (DS:RBX & ~0FFFH) ) or (EPCM(DS:RBX).R = 0) )  
 THEN #PF(DS:RBX);

FI;

(\* Make sure OUTPUTDATA is properly aligned and inside the current enclave \*)

IF ( (DS:RCX is not 16Byte aligned) or (DS:RCX is not within CR\_ELRANGE) )  
 THEN #GP(0); FI;

(\* Make sure DS:RCX is an EPC address and the EPC page is valid \*)

IF ( (DS:RCX does not resolve to an EPC address) or (EPCM(DS:RCX).VALID = 0) )

```

THEN #PF(DS:RCX); FI;

IF (EPCM(DS:RCX).BLOCKED = 1)
  THEN #PF(DS:RCX); FI;

(* Check page parameters for correctness *)
IF ( (EPCM(DS:RCX).PT ≠ PT_REG) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR_ACTIVE_SECS) or (EPCM(DS:RCX).PENDING = 1) or
  (EPCM(DS:RCX).MODIFIED = 1) or (EPCM(DS:RCX).ENCLAVEADDRESS ≠ (DS:RCX & ~0FFFH) ) or (EPCM(DS:RCX).W = 0) )
  THEN #PF(DS:RCX);
FI;

(* Verify RESERVED spaces in KEYREQUEST are valid *)
IF ( (DS:RBX).RESERVED ≠ 0) or (DS:RBX.KEYPOLICY.RESERVED ≠ 0) )
  THEN #GP(0); FI;

TMP_CURRENTSECS ← CR_ACTIVE_SECS;

(* Verify that CONFIGSVN & New Policy bits are not used if KSS is not enabled *)
IF ((TMP_CURRENTSECS.ATTRIBUTES.KSS == 0) AND ((DS:RBX.KEYPOLICY & 0x003C ≠ 0) OR (DS:RBX.CONFIGSVN > 0)))
  THEN #GP(0); FI;

(* Determine which enclave attributes that must be included in the key. Attributes that must always be include INIT & DEBUG *)
REQUIRED_SEALING_MASK[127:0] ← 00000000 00000000 00000000 00000003H;
TMP_ATTRIBUTES ← (DS:RBX.ATTRIBUTEMASK | REQUIRED_SEALING_MASK) & TMP_CURRENTSECS.ATTRIBUTES;

(* Compute MISCSELECT fields to be included *)
TMP_MISCSELECT ← DS:RBX.MISCMASK & TMP_CURRENTSECS.MISCSELECT

CASE (DS:RBX.KEYNAME)
  SEAL_KEY:
    IF (DS:RBX.CPUSVN is beyond current CPU configuration)
      THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_CPUSVN;
        GOTO EXIT;
    FI;
    IF (DS:RBX.ISVSVN > TMP_CURRENTSECS.ISVSVN)
      THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_ISVSVN;
        GOTO EXIT;
    FI;
    IF (DS:RBX.CONFIGSVN > TMP_CURRENTSECS.CONFIGSVN)
      THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_ISVSVN;
        GOTO EXIT;
    FI;

    (*Include enclave identity?*)
    TMP_MRENCLAVE ← 0;
    IF (DS:RBX.KEYPOLICY.MRENCLAVE = 1)
      THEN TMP_MRENCLAVE ← TMP_CURRENTSECS.MRENCLAVE;
    FI;

    (*Include enclave author?*)

```

## SGX INSTRUCTION REFERENCES

```
TMP_MRSIGNER ← 0;
IF (DS:RBX.KEYPOLICY.MRSIGNER = 1)
    THEN TMP_MRSIGNER ← TMP_CURRENTSECS.MRSIGNER;
FI;

(* Include enclave product family ID? *)
TMP_ISVFAMILYID ← 0;
IF (DS:RBX.KEYPOLICY.ISVFAMILYID = 1)
    THEN TMP_ISVFAMILYID ← TMP_CURRENTSECS.ISVFAMILYID;
FI;

(* Include enclave product ID? *)
TMP_ISVPRODID ← 0;
IF (DS:RBX.KEYPOLICY.NOISVPRODID = 0)
    TMP_ISVPRODID ← TMP_CURRENTSECS.ISVPRODID;
FI;

(* Include enclave Config ID? *)
TMP_CONFIGID ← 0;
TMP_CONFIGSVN ← 0;
IF (DS:RBX.KEYPOLICY.CONFIGID = 1)
    TMP_CONFIGID ← TMP_CURRENTSECS.CONFIGID;
    TMP_CONFIGSVN ← DS:RBX.CONFIGSVN;
FI;

(* Include enclave extended product ID? *)
TMP_ISVEXTPRODID ← 0;
IF (DS:RBX.KEYPOLICY.ISVEXTPRODID = 1 )
    TMP_ISVEXTPRODID ← TMP_CURRENTSECS.ISVEXTPRODID;
FI;

//Determine values key is based on
TMP_KEYDEPENDENCIES.KEYNAME ← SEAL_KEY;
TMP_KEYDEPENDENCIES.ISVFAMILYID ← TMP_ISVFAMILYID;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← TMP_ISVEXTPRODID;
TMP_KEYDEPENDENCIES.ISVPRODID ← TMP_ISVPRODID;
TMP_KEYDEPENDENCIES.ISVSVN ← DS:RBX.ISVSVN;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← CR_SGXOWNEREPOCH;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_ATTRIBUTES;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← DS:RBX.ATTRIBUTEMASK;
TMP_KEYDEPENDENCIES.MRENCLAVE ← TMP_MRENCLAVE;
TMP_KEYDEPENDENCIES.MRSIGNER ← TMP_MRSIGNER;
TMP_KEYDEPENDENCIES.KEYID ← DS:RBX.KEYID;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← DS:RBX.CPUSVN;
TMP_KEYDEPENDENCIES.PADDING ← TMP_CURRENTSECS.PADDING;
TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_MISCSELECT;
TMP_KEYDEPENDENCIES.MISCMASK ← ~DS:RBX.MISCMASK;
TMP_KEYDEPENDENCIES.KEYPOLICY ← DS:RBX.KEYPOLICY;
TMP_KEYDEPENDENCIES.CONFIGID ← TMP_CONFIGID;
TMP_KEYDEPENDENCIES.CONFIGSVN ← TMP_CONFIGSVN;
BREAK;
REPORT_KEY:
//Determine values key is based on
TMP_KEYDEPENDENCIES.KEYNAME ← REPORT_KEY;
```

```

TMP_KEYDEPENDENCIES.ISVFAMILYID ← 0;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVSVN ← 0;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← CR_SGXOWNEREPOCH;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_CURRENTSECS.ATTRIBUTES;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← 0;
TMP_KEYDEPENDENCIES.MRENCLAVE ← TMP_CURRENTSECS.MRENCLAVE;
TMP_KEYDEPENDENCIES.MRSIGNER ← 0;
TMP_KEYDEPENDENCIES.KEYID ← DS:RBX.KEYID;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← CR_CPUSVN;
TMP_KEYDEPENDENCIES.PADDING ← HARDCODED_PKCS1_5_PADDING;
TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_CURRENTSECS.MISCSELECT;
TMP_KEYDEPENDENCIES.MISCMASK ← 0;
TMP_KEYDEPENDENCIES.KEYPOLICY ← 0;
TMP_KEYDEPENDENCIES.CONFIGID ← TMP_CURRENTSECS.CONFIGID;
TMP_KEYDEPENDENCIES.CONFIGSVN ← TMP_CURRENTSECS.CONFIGSVN;
BREAK;
EINITTOKEN_KEY:
(* Check ENCLAVE has LAUNCH capability *)
IF (TMP_CURRENTSECS.ATTRIBUTES.LAUNCHKEY = 0)
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_ATTRIBUTE;
        GOTO EXIT;
FI;
IF (DS:RBX.CPUSVN is beyond current CPU configuration)
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_CPUSVN;
        GOTO EXIT;
FI;
IF (DS:RBX.ISVSVN > TMP_CURRENTSECS.ISVSVN)
    THEN
        RFLAGS.ZF ← 1;
        RAX ← SGX_INVALID_ISVSVN;
        GOTO EXIT;
FI;
(* Determine values key is based on *)
TMP_KEYDEPENDENCIES.KEYNAME ← EINITTOKEN_KEY;
TMP_KEYDEPENDENCIES.ISVFAMILYID ← 0;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVPRODID ← TMP_CURRENTSECS.ISVPRODID;
TMP_KEYDEPENDENCIES.ISVSVN ← DS:RBX.ISVSVN;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← CR_SGXOWNEREPOCH;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_ATTRIBUTES;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← 0;
TMP_KEYDEPENDENCIES.MRENCLAVE ← 0;
TMP_KEYDEPENDENCIES.MRSIGNER ← TMP_CURRENTSECS.MRSIGNER;
TMP_KEYDEPENDENCIES.KEYID ← DS:RBX.KEYID;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← DS:RBX.CPUSVN;
TMP_KEYDEPENDENCIES.PADDING ← TMP_CURRENTSECS.PADDING;

```

```

    TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_MISCSELECT;
    TMP_KEYDEPENDENCIES.MISCMASK ← 0;
    TMP_KEYDEPENDENCIES.KEYPOLICY ← 0;
    TMP_KEYDEPENDENCIES.CONFIGID ← 0;
    TMP_KEYDEPENDENCIES.CONFIGSVN ← 0;
    BREAK;
PROVISION_KEY:
(* Check ENCLAVE has PROVISIONING capability *)
    IF (TMP_CURRENTSECS.ATTRIBUTES.PROVISIONKEY = 0)
        THEN
            RFLAGS.ZF ← 1;
            RAX ← SGX_INVALID_ATTRIBUTE;
            GOTO EXIT;
    FI;
    IF (DS:RBX.CPUSVN is beyond current CPU configuration)
        THEN
            RFLAGS.ZF ← 1;
            RAX ← SGX_INVALID_CPUSVN;
            GOTO EXIT;
    FI;
    IF (DS:RBX.ISVSVN > TMP_CURRENTSECS.ISVSVN)
        THEN
            RFLAGS.ZF ← 1;
            RAX ← SGX_INVALID_ISVSVN;
            GOTO EXIT;
    FI;
(* Determine values key is based on *)
    TMP_KEYDEPENDENCIES.KEYNAME ← PROVISION_KEY;
    TMP_KEYDEPENDENCIES.ISVFAMILYID ← 0;
    TMP_KEYDEPENDENCIES.ISVEXTPRODID ← 0;
    TMP_KEYDEPENDENCIES.ISVPRODID ← TMP_CURRENTSECS.ISVPRODID;
    TMP_KEYDEPENDENCIES.ISVSVN ← DS:RBX.ISVSVN;
    TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← 0;
    TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_ATTRIBUTES;
    TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← DS:RBX.ATTRIBUTEMASK;
    TMP_KEYDEPENDENCIES.MRENCLAVE ← 0;
    TMP_KEYDEPENDENCIES.MRSIGNER ← TMP_CURRENTSECS.MRSIGNER;
    TMP_KEYDEPENDENCIES.KEYID ← 0;
    TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← 0;
    TMP_KEYDEPENDENCIES.CPUSVN ← DS:RBX.CPUSVN;
    TMP_KEYDEPENDENCIES.PADDING ← TMP_CURRENTSECS.PADDING;
    TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_MISCSELECT;
    TMP_KEYDEPENDENCIES.MISCMASK ← ~DS:RBX.MISCMASK;
    TMP_KEYDEPENDENCIES.KEYPOLICY ← 0;
    TMP_KEYDEPENDENCIES.CONFIGID ← 0;
    BREAK;
PROVISION_SEAL_KEY:
(* Check ENCLAVE has PROVISIONING capability *)
    IF (TMP_CURRENTSECS.ATTRIBUTES.PROVISIONKEY = 0)
        THEN
            RFLAGS.ZF ← 1;
            RAX ← SGX_INVALID_ATTRIBUTE;
            GOTO EXIT;
    FI;

```



```

IF (DS:RBX.CPUSVN is beyond current CPU configuration)
  THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_CPUSVN;
    GOTO EXIT;
FI;
IF (DS:RBX.ISVSVN > TMP_CURRENTSECS.ISVSVN)
  THEN
    RFLAGS.ZF ← 1;
    RAX ← SGX_INVALID_ISVSVN;
    GOTO EXIT;
FI;
(* Include enclave product family ID? *)
TMP_ISVFAMILYID ← 0;
IF (DS:RBX.KEYPOLICY.ISVFAMILYID = 1)
  THEN TMP_ISVFAMILYID ← TMP_CURRENTSECS.ISVFAMILYID;
FI;

(* Include enclave product ID? *)
TMP_ISVPRODID ← 0;
IF (DS:RBX.KEYPOLICY.NOISVPRODID = 0)
  TMP_ISVPRODID ← TMP_CURRENTSECS.ISVPRODID;
FI;

(* Include enclave Config ID? *)
TMP_CONFIGID ← 0;
TMP_CONFIGSVN ← 0;
IF (DS:RBX.KEYPOLICY.CONFIGID = 1)
  TMP_CONFIGID ← TMP_CURRENTSECS.CONFIGID;
  TMP_CONFIGSVN ← DS:RBX.CONFIGSVN;
FI;

(* Include enclave extended product ID? *)
TMP_ISVEXTPRODID ← 0;
IF (DS:RBX.KEYPOLICY.ISVEXTPRODID = 1)
  TMP_ISVEXTPRODID ← TMP_CURRENTSECS.ISVEXTPRODID;
FI;

(* Determine values key is based on *)
TMP_KEYDEPENDENCIES.KEYNAME ← PROVISION_SEAL_KEY;
TMP_KEYDEPENDENCIES.ISVFAMILYID ← TMP_ISVFAMILYID;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← TMP_ISVEXTPRODID;
TMP_KEYDEPENDENCIES.ISVPRODID ← TMP_ISVPRODID;
TMP_KEYDEPENDENCIES.ISVSVN ← DS:RBX.ISVSVN;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← 0;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← TMP_ATTRIBUTES;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← DS:RBX.ATTRIBUTEMASK;
TMP_KEYDEPENDENCIES.MRENCLAVE ← 0;
TMP_KEYDEPENDENCIES.MRSIGNER ← TMP_CURRENTSECS.MRSIGNER;
TMP_KEYDEPENDENCIES.KEYID ← 0;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← DS:RBX.CPUSVN;
TMP_KEYDEPENDENCIES.PADDING ← TMP_CURRENTSECS.PADDING;
TMP_KEYDEPENDENCIES.MISCSELECT ← TMP_MISCSELECT;

```

## SGX INSTRUCTION REFERENCES

```
TMP_KEYDEPENDENCIES.MISCMASK ← ~DS:RBX.MISCMASK;  
TMP_KEYDEPENDENCIES.KEYPOLICY ← DS:RBX.KEYPOLICY;  
TMP_KEYDEPENDENCIES.CONFIGID ← TMP_CONFIGID;  
TMP_KEYDEPENDENCIES.CONFIGSVN ← TMP_CONFIGSVN;  
BREAK;
```

DEFAULT:

```
(* The value of KEYNAME is invalid *)  
RFLAGS.ZF ← 1;  
RAX ← SGX_INVALID_KEYNAME;  
GOTO EXIT;
```

ESAC;

(\* Calculate the final derived key and output to the address in RCX \*)

```
TMP_OUTPUTKEY ← derivekey(TMP_KEYDEPENDENCIES);  
DS:RCX[15:0] ← TMP_OUTPUTKEY;  
RAX ← 0;  
RFLAGS.ZF ← 0;
```

EXIT:

```
RFLAGS.CF ← 0;  
RFLAGS.PF ← 0;  
RFLAGS.AF ← 0;  
RFLAGS.OF ← 0;  
RFLAGS.SF ← 0;
```

### Flags Affected

ZF is cleared if successful, otherwise ZF is set. CF, PF, AF, OF, SF are cleared.

### Protected Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | If a memory operand effective address is outside the current enclave.<br>If an effective address is not properly aligned.<br>If an effective address is outside the DS segment limit.<br>If KEYREQUEST format is invalid. |
| #PF(error code) | If a page fault occurs in accessing memory.   |

### 64-Bit Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the current enclave.<br>If an effective address is not properly aligned.<br>If an effective address is not canonical.<br>If KEYREQUEST format is invalid. |
| #PF(error code) | If a page fault occurs in accessing memory operands.   |

## EMODPE—Extend an EPC Page Permissions

| Opcode/<br>Instruction     | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 06H<br>ENCLU[EMODPE] | IR    | V/V                          | SGX2                     | This leaf function extends the access rights of an existing EPC page. |

### Instruction Operand Encoding

| Op/En | EAX         | RBX                       | RCX                                      |
|-------|-------------|---------------------------|--|
| IR    | EMODPE (In) | Address of a SECINFO (In) | Address of the destination EPC page (In) |

### Description

This leaf function extends the access rights associated with an existing EPC page in the running enclave. THE RWX bits of the SECINFO parameter are treated as a permissions mask; supplying a value that does not extend the page permissions will have no effect. This instruction leaf can only be executed when inside the enclave.

RBX contains the effective address of a SECINFO structure while RCX contains the effective address of an EPC page. The table below provides additional information on the memory parameter of the EMODPE leaf function.

### EMODPE Memory Parameter Semantics

| SECINFO                              | EPCPAGE                          |
|--------------------------------------|----------------------------------|
| Read access permitted by Non Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

### EMODPE Faulting Conditions

|   |  |
|---|--|
| The operands are not properly aligned.    | If security attributes of the SECINFO page make the page inaccessible.           |
| The EPC page is locked by another thread. | RBX does not contain an effective address in an EPC page in the running enclave. |
| The EPC page is not valid.                | RCX does not contain an effective address of an EPC page in the running enclave. |
| SECINFO contains an invalid request.      |  |

### Concurrency Restrictions

**Table 40-68. Base Concurrency Restrictions of EMODPE**

| Leaf   | Parameter        | Base Concurrency Restrictions |             |                                    |
|--------|------------------|-------------------------------|-------------|------------------------------------|
|        |                  | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EMODPE | Target [DS:RCX]  | Concurrent                    |             |                                    |
|        | SECINFO [DS:RBX] | Concurrent                    |             |                                    |

**Table 40-69. Additional Concurrency Restrictions of EMODPE**

| Leaf   | Parameter        | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------|------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|        |                  | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|        |                  | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EMODPE | Target [DS:RCX]  | Exclusive                                       | #GP         | Concurrent               |             | Concurrent          |             |
|        | SECINFO [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

Operation

Temp Variables in EMODPE Operational Flow

| Name            | Type    | Size (bits) | Description   |
|-----------------|---------|-------------|---|
| SCRATCH_SECINFO | SECINFO | 512         | Scratch storage for holding the contents of DS:RBX. |

IF (DS:RBX is not 64Byte Aligned)  
 THEN #GP(0); FI;

IF (DS:RCX is not 4KByte Aligned)  
 THEN #GP(0); FI;

IF ((DS:RBX is not within CR\_ELRANGE) or (DS:RCX is not within CR\_ELRANGE) )  
 THEN #GP(0); FI;

IF (DS:RBX does not resolve within an EPC)  
 THEN #PF(DS:RBX); FI;

IF (DS:RCX does not resolve within an EPC)  
 THEN #PF(DS:RCX); FI;

IF ( ( EPCM(DS:RBX).VALID = 0) or (EPCM(DS:RBX).R = 0) or (EPCM(DS:RBX).PENDING ≠ 0) or (EPCM(DS:RBX).MODIFIED ≠ 0) or  
 (EPCM(DS:RBX).BLOCKED ≠ 0) or (EPCM(DS:RBX).PT ≠ PT\_REG) or (EPCM(DS:RBX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RBX).ENCLAVEADDRESS ≠ (DS:RBX & ~0xFFF)) )  
 THEN #PF(DS:RBX); FI;

SCRATCH\_SECINFO ← DS:RBX;

(\* Check for mis-configured SECINFO flags\*)  
 IF (SCRATCH\_SECINFO reserved fields are not zero )  
 THEN #GP(0); FI;

(\* Check security attributes of the EPC page \*)  
 IF ( ( EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PENDING ≠ 0) or (EPCM(DS:RCX).MODIFIED ≠ 0) or  
 (EPCM(DS:RCX).BLOCKED ≠ 0) or (EPCM(DS:RCX).PT ≠ PT\_REG) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) )  
 THEN #PF(DS:RCX); FI;

(\* Check the EPC page for concurrency \*)  
 IF (EPC page in use by another SGX2 instruction)  
 THEN #GP(0); FI;

(\* Re-Check security attributes of the EPC page \*)  
 IF ( ( EPCM(DS:RCX).VALID = 0) or (EPCM(DS:RCX).PENDING ≠ 0) or (EPCM(DS:RCX).MODIFIED ≠ 0) or  
 (EPCM(DS:RCX).PT ≠ PT\_REG) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or  
 (EPCM(DS:RCX).ENCLAVEADDRESS ≠ DS:RCX))  
 THEN #PF(DS:RCX); FI;

(\* Check for mis-configured SECINFO flags\*)  
 IF ( ( EPCM(DS:RCX).R = 0) and (SCRATCH\_SECINFO.FLAGS.R = 0) and (SCRATCH\_SECINFO.FLAGS.W ≠ 0) ) )  
 THEN #GP(0); FI;

(\* Update EPCM permissions \*)

EPCM(DS:RCX).R ← EPCM(DS:RCX).R | SCRATCH\_SECINFO.FLAGS.R;  
 EPCM(DS:RCX).W ← EPCM(DS:RCX).W | SCRATCH\_SECINFO.FLAGS.W;  
 EPCM(DS:RCX).X ← EPCM(DS:RCX).X | SCRATCH\_SECINFO.FLAGS.X;

### Flags Affected

None

### Protected Mode Exceptions

#GP(0)                If a memory operand effective address is outside the DS segment limit.  
                           If a memory operand is not properly aligned.  
                           If a memory operand is locked.  
 #PF(error code)    If a page fault occurs in accessing memory operands.

### 64-Bit Mode Exceptions

#GP(0)                If a memory operand is non-canonical form.  
                           If a memory operand is not properly aligned.  
                           If a memory operand is locked.  
 #PF(error code)    If a page fault occurs in accessing memory operands.

## EReport—Create a Cryptographic Report of the Enclave

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 00H<br>ENCLU[EReport] | IR    | V/V                          | SGX1                     | This leaf function creates a cryptographic report of the enclave. |

### Instruction Operand Encoding

| Op/En | EAX          | RBX                        | RCX                        | RDX  |
|-------|--------------|----------------------------|----------------------------|--|
| IR    | EReport (In) | Address of TARGETINFO (In) | Address of REPORTDATA (In) | Address where the REPORT is written to in an OUTPUTDATA (In) |

### Description

This leaf function creates a cryptographic REPORT that describes the contents of the enclave. This instruction leaf can only be executed when inside the enclave. The cryptographic report can be used by other enclaves to determine that the enclave is running on the same platform.

RBX contains the effective address of the MRENCLAVE value of the enclave that will authenticate the REPORT output, using the REPORT key delivered by EGETKEY command for that enclave. RCX contains the effective address of a 64-byte REPORTDATA structure, which allows the caller of the instruction to associate data with the enclave from which the instruction is called. RDX contains the address where the REPORT will be output by the instruction.

### EReport Memory Parameter Semantics

| TARGETINFO             | REPORTDATA             | OUTPUTDATA                   |
|------------------------|------------------------|------------------------------|
| Read access by Enclave | Read access by Enclave | Read/Write access by Enclave |

This instruction leaf perform the following:

1. Validate the 3 operands (RBX, RCX, RDX) are inside the enclave.
2. Compute a report key for the target enclave, as indicated by the value located in RBX(TARGETINFO).
3. Assemble the enclave SECS data to complete the REPORT structure (including the data provided using the RCX (REPORTDATA) operand).
4. Computes a cryptographic hash over REPORT structure.
5. Add the computed hash to the REPORT structure.
6. Output the completed REPORT structure to the address in RDX (OUTPUTDATA).

The instruction fails if the operands are not properly aligned.

CR\_REPORT\_KEYID, used to provide key wearout protection, is populated with a statistically unique value on boot of the platform by a trusted entity within the SGX TCB.

The instruction faults if any of the following:

### EReport Faulting Conditions

|  |  |
|--|--|
| An effective address not properly aligned. | An memory address does not resolve in an EPC page. |
| If accessing an invalid EPC page.          | If the EPC page is blocked.                        |
| May page fault.                            |  |

## Concurrency Restrictions

Table 40-70. Base Concurrency Restrictions of EREPORT

| Leaf    | Parameter           | Base Concurrency Restrictions |             |                                    |
|---------|---------------------|-------------------------------|-------------|------------------------------------|
|         |                     | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| EREPORT | TARGETINFO [DS:RBX] | Concurrent                    |             |                                    |
|         | REPORTDATA [DS:RCX] | Concurrent                    |             |                                    |
|         | OUTPUTDATA [DS:RDX] | Concurrent                    |             |                                    |

Table 40-71. Additional Concurrency Restrictions of EREPORT

| Leaf    | Parameter           | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|---------------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |                     | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |                     | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EREPORT | TARGETINFO [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|         | REPORTDATA [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|         | OUTPUTDATA [DS:RDX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

## Operation

## Temp Variables in EREPORT Operational Flow

| Name                | Type | Size (bits) | Description  |
|---------------------|------|-------------|--|
| TMP_ATTRIBUTES      |      | 32          | Physical address of SECS of the enclave to which source operand belongs. |
| TMP_CURRENTSECS     |      |             | Address of the SECS for the currently executing enclave.                 |
| TMP_KEYDEPENDENCIES |      |             | Temp space for key derivation.   |
| TMP_REPORTKEY       |      | 128         | REPORTKEY generated by the instruction.                                  |
| TMP_REPORT          |      | 3712        |  |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

(\* Address verification for TARGETINFO (RBX) \*)

IF ( (DS:RBX is not 512Byte Aligned) or (DS:RBX is not within CR\_ELRange) )  
THEN #GP(0); FI;

IF (DS:RBX does not resolve within an EPC)  
THEN #PF(DS:RBX); FI;

IF (EPCM(DS:RBX).VALID = 0)  
THEN #PF(DS:RBX); FI;

IF (EPCM(DS:RBX).BLOCKED = 1)  
THEN #PF(DS:RBX); FI;

(\* Check page parameters for correctness \*)

IF ( (EPCM(DS:RBX).PT ≠ PT\_REG) or (EPCM(DS:RBX).ENCLAVESECS ≠ CR\_ACTIVE\_SECS) or (EPCM(DS:RBX).PENDING = 1) or  
(EPCM(DS:RBX).MODIFIED = 1) or (EPCM(DS:RBX).ENCLAVEADDRESS ≠ (DS:RBX & ~0FFFH) ) or (EPCM(DS:RBX).R = 0) )

## SGX INSTRUCTION REFERENCES

```
THEN #PF(DS:RBX);  
FI;
```

```
(* Address verification for REPORTDATA (RCX) *)  
IF ( (DS:RCX is not 128Byte Aligned) or (DS:RCX is not within CR_ELRANGE) )  
    THEN #GP(0); FI;
```

```
IF (DS:RCX does not resolve within an EPC)  
    THEN #P(DS:RCX); FI;
```

```
IF (EPCM(DS:RCX).VALID = 0)  
    THEN #PF(DS:RCX); FI;
```

```
IF (EPCM(DS:RCX).BLOCKED = 1) )  
    THEN #PF(DS:RCX); FI;
```

```
(* Check page parameters for correctness *)  
IF ( (EPCM(DS:RCX).PT ≠ PT_REG) or (EPCM(DS:RCX).ENCLAVESECS ≠ CR_ACTIVE_SECS) or (EPCM(DS:RCX).PENDING = 1) or  
    (EPCM(DS:RCX).MODIFIED = 1) or (EPCM(DS:RCX).ENCLAVEADDRESS ≠ (DS:RCX & ~OFFFH) ) or (EPCM(DS:RCX).R = 0) )  
    THEN #PF(DS:RCX);  
FI;
```

```
(* Address verification for OUTPUTDATA (RDX) *)  
IF ( (DS:RDX is not 512Byte Aligned) or (DS:RDX is not within CR_ELRANGE) )  
    THEN #GP(0); FI;
```

```
IF (DS:RDX does not resolve within an EPC)  
    THEN #PF(DS:RDX); FI;
```

```
IF (EPCM(DS:RDX).VALID = 0)  
    THEN #PF(DS:RDX); FI;
```

```
IF (EPCM(DS:RDX).BLOCKED = 1) )  
    THEN #PF(DS:RDX); FI;
```

```
(* Check page parameters for correctness *)  
IF ( (EPCM(DS:RDX).PT ≠ PT_REG) or (EPCM(DS:RDX).ENCLAVESECS ≠ CR_ACTIVE_SECS) or (EPCM(DS:RCX).PENDING = 1) or  
    (EPCM(DS:RCX).MODIFIED = 1) or (EPCM(DS:RDX).ENCLAVEADDRESS ≠ (DS:RDX & ~OFFFH) ) or (EPCM(DS:RDX).W = 0) )  
    THEN #PF(DS:RDX);  
FI;
```

```
(* REPORT MAC needs to be computed over data which cannot be modified *)  
TMP_REPORT.CPUSVN ← CR_CPUSVN;  
TMP_REPORT.ISVFAMILYID ← TMP_CURRENTSECS.ISVFAMILYID;  
TMP_REPORT.ISVEXTPRODID ← TMP_CURRENTSECS.ISVEXTPRODID;  
TMP_REPORT.ISVPRODID ← TMP_CURRENTSECS.ISVPRODID;  
TMP_REPORT.ISVSVN ← TMP_CURRENTSECS.ISVSVN;  
TMP_REPORT.ATTRIBUTES ← TMP_CURRENTSECS.ATTRIBUTES;  
TMP_REPORT.REPORTDATA ← DS:RCX[511:0];  
TMP_REPORT.MRENCLAVE ← TMP_CURRENTSECS.MRENCLAVE;  
TMP_REPORT.MRSIGNER ← TMP_CURRENTSECS.MRSIGNER;  
TMP_REPORT.MRRESERVED ← 0;  
TMP_REPORT.KEYID[255:0] ← CR_REPORT_KEYID;  
TMP_REPORT.MISCSELECT ← TMP_CURRENTSECS.MISCSELECT;
```



```
TMP_REPORT.CONFIGID ← TMP_CURRENTSECS.CONFIGID;
TMP_REPORT.CONFIGSVN ← TMP_CURRENTSECS.CONFIGSVN;
```

(\* Derive the report key \*)

```
TMP_KEYDEPENDENCIES.KEYNAME ← REPORT_KEY;
TMP_KEYDEPENDENCIES.ISVFAMILYID ← 0;
TMP_KEYDEPENDENCIES.ISVEXTPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVPRODID ← 0;
TMP_KEYDEPENDENCIES.ISVSVN ← 0;
TMP_KEYDEPENDENCIES.SGXOWNEREPOCH ← CR_SGXOWNEREPOCH;
TMP_KEYDEPENDENCIES.ATTRIBUTES ← DS:RBX.ATTRIBUTES;
TMP_KEYDEPENDENCIES.ATTRIBUTESMASK ← 0;
TMP_KEYDEPENDENCIES.MRENCLAVE ← DS:RBX.MEASUREMENT;
TMP_KEYDEPENDENCIES.MRSIGNER ← 0;
TMP_KEYDEPENDENCIES.KEYID ← TMP_REPORT.KEYID;
TMP_KEYDEPENDENCIES.SEAL_KEY_FUSES ← CR_SEAL_FUSES;
TMP_KEYDEPENDENCIES.CPUSVN ← CR_CPUSVN;
TMP_KEYDEPENDENCIES.PADDING ← TMP_CURRENTSECS.PADDING;
TMP_KEYDEPENDENCIES.MISCSELECT ← DS:RBX.MISCSELECT;
TMP_KEYDEPENDENCIES.MISCMASK ← 0;
TMP_KEYDEPENDENCIES.KEYPOLICY ← 0;
TMP_KEYDEPENDENCIES.CONFIGID ← DS:RBX.CONFIGID;
TMP_KEYDEPENDENCIES.CONFIGSVN ← DS:RBX.CONFIGSVN;
```

(\* Calculate the derived key\*)

```
TMP_REPORTKEY ← derive_key(TMP_KEYDEPENDENCIES);
```

(\* call cryptographic CMAC function, CMAC data are not including MAC&KEYID \*)

```
TMP_REPORT.MAC ← cmac(TMP_REPORTKEY, TMP_REPORT[3071:0]);
DS:RDX[3455:0] ← TMP_REPORT;
```

### Flags Affected

None

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If the address in RCS is outside the DS segment limit.<br>If a memory operand is not properly aligned.<br>If a memory operand is not in the current enclave. |
| #PF(error code) | If a page fault occurs in accessing memory operands.   |

### 64-Bit Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | If RCX is non-canonical form.<br>If a memory operand is not properly aligned.<br>If a memory operand is not in the current enclave. |
| #PF(error code) | If a page fault occurs in accessing memory operands.  |

## ERESUME—Re-Enters an Enclave

| Opcode/<br>Instruction      | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|-----------------------------|-------|------------------------------|--------------------------|---|
| EAX = 03H<br>ENCLU[ERESUME] | IR    | V/V                          | SGX1                     | This leaf function is used to re-enter an enclave after an interrupt. |

### Instruction Operand Encoding

| Op/En | RAX          | RBX                   | RCX                 |
|-------|--------------|-----------------------|---------------------|
| IR    | ERESUME (In) | Address of a TCS (In) | Address of AEP (In) |

### Description

The ENCLU[ERESUME] instruction resumes execution of an enclave that was interrupted due to an exception or interrupt, using the machine state previously stored in the SSA.

### ERESUME Memory Parameter Semantics

| TCS                       |
|---------------------------|
| Enclave read/write access |

The instruction faults if any of the following:

|  |   |
|--|---|
| Address in RBX is not properly aligned.                    | Any TCS.FLAGS's must-be-zero bit is not zero.   |
| TCS pointed to by RBX is not valid or available or locked. | Current 32/64 mode does not match the enclave mode in SECS.ATTRIBUTES.MODE64.                   |
| The SECS is in use by another enclave.                     | Either of TCS-specified FS and GS segment is not a subset of the current DS segment.            |
| Any one of DS, ES, CS, SS is not zero.                     | If XSAVE available, CR4.OSXSAVE = 0, but SECS.ATTRIBUTES.XFRM ≠ 3.                              |
| CR4.OSFXSR ≠ 1.  | If CR4.OSXSAVE = 1, SECS.ATTRIBUTES.XFRM is not a subset of XCR0.                               |
| Offsets 520-535 of the XSAVE area not 0.                   | The bit vector stored at offset 512 of the XSAVE area must be a subset of SECS.ATTRIBUTES.XFRM. |
| The SSA frame is not valid or in use.                      |   |

The following operations are performed by ERESUME:

- RSP and RBP are saved in the current SSA frame on EENTER and are automatically restored on EEXIT or an asynchronous exit due to any Interrupt event.
- The AEP contained in RCX is stored into the TCS for use by AEXs. FS and GS (including hidden portions) are saved and new values are constructed using TCS.OFSBASE/GSBASE (32 and 64-bit mode) and TCS.OFSLIMIT/GSLIMIT (32-bit mode only). The resulting segments must be a subset of the DS segment.
- If CR4.OSXSAVE == 1, XCR0 is saved and replaced by SECS.ATTRIBUTES.XFRM. The effect of RFLAGS.TF depends on whether the enclave entry is opt-in or opt-out (see Section 42.1.2):
  - On opt-out entry, TF is saved and cleared (it is restored on EEXIT or AEX). Any attempt to set TF via a POPF instruction while inside the enclave clears TF (see Section 42.2.5).
  - On opt-in entry, a single-step debug exception is pending on the instruction boundary immediately after EENTER (see Section 42.2.3).
- All code breakpoints that do not overlap with ELRANGE are also suppressed. If the entry is an opt-out entry, all code and data breakpoints that overlap with the ELRANGE are suppressed.

- On opt-out entry, a number of performance monitoring counters and behaviors are modified or suppressed (see Section 42.2.3):
  - All performance monitoring activity on the current thread is suppressed except for incrementing and firing of FIXED\_CTR1 and FIXED\_CTR2.
  - PEBS is suppressed.
  - AnyThread counting on other threads is demoted to MyThread mode and IA32\_PERF\_GLOBAL\_STATUS[60] on that thread is set.
  - If the opt-out entry on a hardware thread results in suppression of any performance monitoring, then the processor sets IA32\_PERF\_GLOBAL\_STATUS[60] and IA32\_PERF\_GLOBAL\_STATUS[63].

## Concurrency Restrictions

**Table 40-72. Base Concurrency Restrictions of ERESUME**

| Leaf    | Parameter    | Base Concurrency Restrictions |             |                                    |
|---------|--------------|-------------------------------|-------------|------------------------------------|
|         |              | Access                        | On Conflict | SGX_CONFLICT VM Exit Qualification |
| ERESUME | TCS [DS:RBX] | Shared                        | #GP         |                                    |

**Table 40-73. Additional Concurrency Restrictions of ERESUME**

| Leaf    | Parameter    | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------|--------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|         |              | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|         |              | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| ERESUME | TCS [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

## Operation

### Temp Variables in ERESUME Operational Flow

| Name              | Type              | Size  | Description  |
|-------------------|-------------------|-------|--|
| TMP_FSBASE        | Effective Address | 32/64 | Proposed base address for FS segment.  |
| TMP_GSBASE        | Effective Address | 32/64 | Proposed base address for GS segment.  |
| TMP_FSLIMIT       | Effective Address | 32/64 | Highest legal address in proposed FS segment.                                |
| TMP_GSLIMIT       | Effective Address | 32/64 | Highest legal address in proposed GS segment.                                |
| TMP_TARGET        | Effective Address | 32/64 | Address of first instruction inside enclave at which execution is to resume. |
| TMP_SECS          | Effective Address | 32/64 | Physical address of SECS for this enclave.                                   |
| TMP_SSA           | Effective Address | 32/64 | Address of current SSA frame.  |
| TMP_XSIZE         | integer           | 64    | Size of XSAVE area based on SECS.ATTRIBUTES.XFRM.                            |
| TMP_SSA_PAGE      | Effective Address | 32/64 | Pointer used to iterate over the SSA pages in the current frame.             |
| TMP_GPR           | Effective Address | 32/64 | Address of the GPR area within the current SSA frame.                        |
| TMP_BRANCH_RECORD | LBR Record        |       | From/to addresses to be pushed onto the LBR stack.                           |

TMP\_MODE64 ← ((IA32\_EFER.LMA = 1) && (CS.L = 1));

(\* Make sure DS is usable, expand up \*)

IF (TMP\_MODE64 = 0 and (DS not usable or ((DS[S] = 1) and (DS[bit 11] = 0) and DS[bit 10] = 1)))  
 THEN #GP(0); FI;

## SGX INSTRUCTION REFERENCES

(\* Check that CS, SS, DS, ES.base is 0 \*)

```
IF (TMP_MODE64 = 0)
  THEN
    IF(CS.base ≠ 0 or DS.base ≠ 0) #GP(0); FI;
    IF(ES usable and ES.base ≠ 0) #GP(0); FI;
    IF(SS usable and SS.base ≠ 0) #GP(0); FI;
    IF(SS usable and SS.B = 0) #GP(0); FI;
FI;
```

```
IF (DS:RBX is not 4KByte Aligned)
  THEN #GP(0); FI;
```

```
IF (DS:RBX does not resolve within an EPC)
  THEN #PF(DS:RBX); FI;
```

(\* Check AEP is canonical\*)

```
IF (TMP_MODE64 = 1 and (CS:RCX is not canonical) )
  THEN #GP(0); FI;
```

(\* Check concurrency of TCS operation\*)

```
IF (Other Intel SGX instructions is operating on TCS)
  THEN #GP(0); FI;
```

(\* TCS verification \*)

```
IF (EPCM(DS:RBX).VALID = 0)
  THEN #PF(DS:RBX); FI;
```

```
IF (EPCM(DS:RBX).BLOCKED = 1)
  THEN #PF(DS:RBX); FI;
```

```
IF ((EPCM(DS:RBX).PENDING = 1) or (EPCM(DS:RBX).MODIFIED = 1))
  THEN #PF(DS:RBX); FI;
```

```
IF ( (EPCM(DS:RBX).ENCLAVEADDRESS ≠ DS:RBX) or (EPCM(DS:RBX).PT ≠ PT_TCS) )
  THEN #PF(DS:RBX); FI;
```

```
IF ( (DS:RBX).OSSA is not 4KByte Aligned)
  THEN #GP(0); FI;
```

(\* Check proposed FS and GS \*)

```
IF ( ( (DS:RBX).OFSBASE is not 4KByte Aligned) or ( (DS:RBX).OGSBASE is not 4KByte Aligned) )
  THEN #GP(0); FI;
```

(\* Get the SECS for the enclave in which the TCS resides \*)

```
TMP_SECS ← Address of SECS for TCS;
```

(\* Make sure that the FLAGS field in the TCS does not have any reserved bits set \*)

```
IF ( ( (DS:RBX).FLAGS & & FFFFFFFF) ≠ 0)
  THEN #GP(0); FI;
```

(\* SECS must exist and enclave must have previously been EINITted \*)

```
IF (the enclave is not already initialized)
  THEN #GP(0); FI;
```

(\* make sure the logical processor's operating mode matches the enclave \*)

```
IF ( (TMP_MODE64 ≠ TMP_SECS.ATTRIBUTES.MODE64BIT) )
  THEN #GP(0); FI;
```

```
IF (CR4.OSFXSR = 0)
  THEN #GP(0); FI;
```

(\* Check for legal values of SECS.ATTRIBUTES.XFRM \*)

```
IF (CR4.OSXSAVE = 0)
  THEN
    IF (TMP_SECS.ATTRIBUTES.XFRM ≠ 03H) THEN #GP(0); FI;
  ELSE
    IF ( (TMP_SECS.ATTRIBUTES.XFRM & XCRO) ≠ TMP_SECS.ATTRIBUTES.XFRM) THEN #GP(0); FI;
  FI;
```

(\* Make sure the SSA contains at least one active frame \*)

```
IF ( (DS:RBX).CSSA = 0)
  THEN #GP(0); FI;
```

(\* Compute linear address of SSA frame \*)

```
TMP_SSA ← (DS:RBX).OSSA + TMP_SECS.BASEADDR + 4096 * TMP_SECS.SSAFRAMESIZE * ( (DS:RBX).CSSA - 1);
TMP_XSIZE ← compute_XSAVE_frame_size(TMP_SECS.ATTRIBUTES.XFRM);
```

```
FOR EACH TMP_SSA_PAGE = TMP_SSA to TMP_SSA + TMP_XSIZE
```

```
  (* Check page is read/write accessible *)
```

```
  Check that DS:TMP_SSA_PAGE is read/write accessible;
```

```
  If a fault occurs, release locks, abort and deliver that fault;
```

```
  IF (DS:TMP_SSA_PAGE does not resolve to EPC page)
```

```
    THEN #PF(DS:TMP_SSA_PAGE); FI;
```

```
  IF (EPCM(DS:TMP_SSA_PAGE).VALID = 0)
```

```
    THEN #PF(DS:TMP_SSA_PAGE); FI;
```

```
  IF (EPCM(DS:TMP_SSA_PAGE).BLOCKED = 1)
```

```
    THEN #PF(DS:TMP_SSA_PAGE); FI;
```

```
  IF ((EPCM(DS:TMP_SSA_PAGE).PENDING = 1) or (EPCM(DS:TMP_SSA_PAGE).MODIFIED = 1))
```

```
    THEN #PF(DS:TMP_SSA_PAGE); FI;
```

```
  IF ( ( EPCM(DS:TMP_SSA_PAGE).ENCLAVEADDRESS ≠ DS:TMP_SSA_PAGE) or (EPCM(DS:TMP_SSA_PAGE).PT ≠ PT_REG) or
```

```
    (EPCM(DS:TMP_SSA_PAGE).ENCLAVESECS ≠ EPCM(DS:RBX).ENCLAVESECS) or
```

```
    (EPCM(DS:TMP_SSA_PAGE).R = 0) or (EPCM(DS:TMP_SSA_PAGE).W = 0) )
```

```
    THEN #PF(DS:TMP_SSA_PAGE); FI;
```

```
  CR_XSAVE_PAGE_n ← Physical_Address(DS:TMP_SSA_PAGE);
```

```
ENDFOR
```

(\* Compute address of GPR area\*)

```
TMP_GPR ← TMP_SSA + 4096 * DS:TMP_SECS.SSAFRAMESIZE -- sizeof(GPRSGX_AREA);
```

```
Check that DS:TMP_SSA_PAGE is read/write accessible;
```

```
If a fault occurs, release locks, abort and deliver that fault;
```

```
IF (DS:TMP_GPR does not resolve to EPC page)
```

```
  THEN #PF(DS:TMP_GPR); FI;
```

```
IF (EPCM(DS:TMP_GPR).VALID = 0)
```

```
  THEN #PF(DS:TMP_GPR); FI;
```

```
IF (EPCM(DS:TMP_GPR).BLOCKED = 1)
```

```
  THEN #PF(DS:TMP_GPR); FI;
```

```
IF ((EPCM(DS:TMP_GPR).PENDING = 1) or (EPCM(DS:TMP_GPR).MODIFIED = 1))
```

```
  THEN #PF(DS:TMP_GPR); FI;
```

## SGX INSTRUCTION REFERENCES

```
IF ( ( EPCM(DS:TMP_GPR).ENCLAVEADDRESS ≠ DS:TMP_GPR) or (EPCM(DS:TMP_GPR).PT ≠ PT_REG) or
(EPCM(DS:TMP_GPR).ENCLAVESECS ≠ EPCM(DS:RBX).ENCLAVESECS) or
(EPCM(DS:TMP_GPR).R = 0) or (EPCM(DS:TMP_GPR).W = 0) )
THEN #PF(DS:TMP_GPR); FI;
```

```
IF (TMP_MODE64 = 0)
THEN
    IF (TMP_GPR + (GPR_SIZE - 1) is not in DS segment) THEN #GP(0); FI;
FI;
```

```
CR_GPR_PA ← Physical_Address (DS: TMP_GPR);
```

```
TMP_TARGET ← (DS:TMP_GPR).RIP;
IF (TMP_MODE64 = 1)
THEN
    IF (TMP_TARGET is not canonical) THEN #GP(0); FI;
ELSE
    IF (TMP_TARGET > CS limit) THEN #GP(0); FI;
FI;
```

(\* Check proposed FS/GS segments fall within DS \*)

```
IF (TMP_MODE64 = 0)
THEN
    TMP_FSBASE ← (DS:RBX).OFSBASE + TMP_SECS.BASEADDR;
    TMP_FSLIMIT ← (DS:RBX).OFSBASE + TMP_SECS.BASEADDR + (DS:RBX).FSLIMIT;
    TMP_GSBASE ← (DS:RBX).OGSBASE + TMP_SECS.BASEADDR;
    TMP_GSLIMIT ← (DS:RBX).OGSBASE + TMP_SECS.BASEADDR + (DS:RBX).GSLIMIT;
    (* if FS wrap-around, make sure DS has no holes*)
    IF (TMP_FSLIMIT < TMP_FSBASE)
    THEN
        IF (DS.limit < 4GB) THEN #GP(0); FI;
    ELSE
        IF (TMP_FSLIMIT > DS.limit) THEN #GP(0); FI;
    FI;
    (* if GS wrap-around, make sure DS has no holes*)
    IF (TMP_GSLIMIT < TMP_GSBASE)
    THEN
        IF (DS.limit < 4GB) THEN #GP(0); FI;
    ELSE
        IF (TMP_GSLIMIT > DS.limit) THEN #GP(0); FI;
    FI;
ELSE
    TMP_FSBASE ← DS:TMP_GPR.FSBASE;
    TMP_GSBASE ← DS:TMP_GPR.GSBASE;
    IF ( (TMP_FSBASE is not canonical) or (TMP_GSBASE is not canonical))
    THEN #GP(0); FI;
FI;
```

(\* Ensure the enclave is not already active and this thread is the only one using the TCS\*)

```
IF (DS:RBX.STATE = ACTIVE)
THEN #GP(0); FI;
```

(\* SECS.ATTRIBUTES.XFRM selects the features to be saved. \*)

(\* CR\_XSAVE\_PAGE\_n: A list of 1 or more physical address of pages that contain the XSAVE area. \*)

```

XRSTOR(TMP_MODE64, SECS.ATTRIBUTES.XFRM, CR_XSAVE_PAGE_n);

IF (XRSTOR failed with #GP)
  THEN
    DS:RBX.STATE ← INACTIVE;
    #GP(0);
FI;

CR_ENCLAVE_MODE ← 1;
CR_ACTIVE_SECS ← TMP_SECS;
CR_ELRRANGE ← (TMP_SECS.BASEADDR, TMP_SECS.SIZE);

(* Save state for possible AEXs *)
CR_TCS_PA ← Physical_Address (DS:RBX);
CR_TCS_LA ← RBX;
CR_TCS_LA.AEP ← RCX;

(* Save the hidden portions of FS and GS *)
CR_SAVE_FS_selector ← FS.selector;
CR_SAVE_FS_base ← FS.base;
CR_SAVE_FS_limit ← FS.limit;
CR_SAVE_FS_access_rights ← FS.access_rights;
CR_SAVE_GS_selector ← GS.selector;
CR_SAVE_GS_base ← GS.base;
CR_SAVE_GS_limit ← GS.limit;
CR_SAVE_GS_access_rights ← GS.access_rights;

RIP ← TMP_TARGET;

Restore_GPRs from DS:TMP_GPR;

(*Restore the RFLAGS values from SSA*)
RFLAGS.CF ← DS:TMP_GPR.RFLAGS.CF;
RFLAGS.PF ← DS:TMP_GPR.RFLAGS.PF;
RFLAGS.AF ← DS:TMP_GPR.RFLAGS.AF;
RFLAGS.ZF ← DS:TMP_GPR.RFLAGS.ZF;
RFLAGS.SF ← DS:TMP_GPR.RFLAGS.SF;
RFLAGS.DF ← DS:TMP_GPR.RFLAGS.DF;
RFLAGS.OF ← DS:TMP_GPR.RFLAGS.OF;
RFLAGS.NT ← DS:TMP_GPR.RFLAGS.NT;
RFLAGS.AC ← DS:TMP_GPR.RFLAGS.AC;
RFLAGS.ID ← DS:TMP_GPR.RFLAGS.ID;
RFLAGS.RF ← DS:TMP_GPR.RFLAGS.RF;
RFLAGS.VM ← 0;
IF (RFLAGS.IOPL = 3)
  THEN RFLAGS.IF = DS:TMP_GPR.IF; FI;

IF (TCS.FLAGS.OPTIN = 0)
  THEN RFLAGS.TF = 0; FI;

(* If XSAVE is enabled, save XCRO and replace it with SECS.ATTRIBUTES.XFRM*)
IF (CR4.OSXSAVE = 1)
  CR_SAVE_XCRO ← XCRO;
  XCRO ← TMP_SECS.ATTRIBUTES.XFRM;

```

## SGX INSTRUCTION REFERENCES

FI;

(\* Pop the SSA stack\*)  
(DS:RBX).CSSA  $\leftarrow$  (DS:RBX).CSSA - 1;

(\* Do the FS/GS swap \*)  
FS.base  $\leftarrow$  TMP\_FSBASE;  
FS.limit  $\leftarrow$  DS:RBX.FSLIMIT;  
FS.type  $\leftarrow$  0001b;  
FS.W  $\leftarrow$  DS.W;  
FS.S  $\leftarrow$  1;  
FS.DPL  $\leftarrow$  DS.DPL;  
FS.G  $\leftarrow$  1;  
FS.B  $\leftarrow$  1;  
FS.P  $\leftarrow$  1;  
FS.AVL  $\leftarrow$  DS.AVL;  
FS.L  $\leftarrow$  DS.L;  
FS.unusable  $\leftarrow$  0;  
FS.selector  $\leftarrow$  0BH;

GS.base  $\leftarrow$  TMP\_GSBASE;  
GS.limit  $\leftarrow$  DS:RBX.GSLIMIT;  
GS.type  $\leftarrow$  0001b;  
GS.W  $\leftarrow$  DS.W;  
GS.S  $\leftarrow$  1;  
GS.DPL  $\leftarrow$  DS.DPL;  
GS.G  $\leftarrow$  1;  
GS.B  $\leftarrow$  1;  
GS.P  $\leftarrow$  1;  
GS.AVL  $\leftarrow$  DS.AVL;  
GS.L  $\leftarrow$  DS.L;  
GS.unusable  $\leftarrow$  0;  
GS.selector  $\leftarrow$  0BH;

CR\_DBGOPTIN  $\leftarrow$  TCS.FLAGS.DBGOPTIN;  
Suppress all code breakpoints that are outside ELRANGE;

IF (CR\_DBGOPTIN = 0)  
  THEN  
    Suppress all code breakpoints that overlap with ELRANGE;  
    CR\_SAVE\_TF  $\leftarrow$  RFLAGS.TF;  
    RFLAGS.TF  $\leftarrow$  0;  
    Suppress any MTF VM exits during execution of the enclave;  
    Clear all pending debug exceptions;  
    Clear any pending MTF VM exit;  
  ELSE  
    Clear all pending debug exceptions;  
    Clear pending MTF VM exits;  
FI;

(\* Assure consistent translations \*)  
Flush\_linear\_context;  
Clear\_Monitor\_FSM;  
Allow\_front\_end\_to\_begin\_fetch\_at\_new\_RIP;



**Flags Affected**

RFLAGS.TF is cleared on opt-out entry

**Protected Mode Exceptions**

|                 |   |
|-----------------|---|
| #GP(0)          | <p>If DS:RBX is not page aligned.</p> <p>If the enclave is not initialized.</p> <p>If the thread is not in the INACTIVE state.</p> <p>If CS, DS, ES or SS bases are not all zero.</p> <p>If executed in enclave mode.</p> <p>If part or all of the FS or GS segment specified by TCS is outside the DS segment.</p> <p>If any reserved field in the TCS FLAG is set.</p> <p>If the target address is not within the CS segment.</p> <p>If CR4.OSFXSR = 0.</p> <p>If CR4.OSXSAVE = 0 and SECS.ATTRIBUTES.XFRM ≠ 3.</p> <p>If CR4.OSXSAVE = 1 and SECS.ATTRIBUTES.XFRM is not a subset of XCR0.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory.</p> <p>If DS:RBX does not point to a valid TCS.</p> <p>If one or more pages of the current SSA frame are not readable/writable, or do not resolve to a valid PT_REG EPC page.</p>  |

**64-Bit Mode Exceptions**

|                 |   |
|-----------------|---|
| #GP(0)          | <p>If DS:RBX is not page aligned.</p> <p>If the enclave is not initialized.</p> <p>If the thread is not in the INACTIVE state.</p> <p>If CS, DS, ES or SS bases are not all zero.</p> <p>If executed in enclave mode.</p> <p>If part or all of the FS or GS segment specified by TCS is outside the DS segment.</p> <p>If any reserved field in the TCS FLAG is set.</p> <p>If the target address is not canonical.</p> <p>If CR4.OSFXSR = 0.</p> <p>If CR4.OSXSAVE = 0 and SECS.ATTRIBUTES.XFRM ≠ 3.</p> <p>If CR4.OSXSAVE = 1 and SECS.ATTRIBUTES.XFRM is not a subset of XCR0.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory operands.</p> <p>If DS:RBX does not point to a valid TCS.</p> <p>If one or more pages of the current SSA frame are not readable/writable, or do not resolve to a valid PT_REG EPC page.</p>   |

## 40.5 INTEL® SGX VIRTUALIZATION LEAF FUNCTION REFERENCE

Leaf functions available with the ENCLV instruction mnemonic are covered in this section. In general, each instruction leaf requires EAX to specify the leaf function index and/or additional implicit registers specifying leaf-specific input parameters. An instruction operand encoding table provides details of each implicit register usage and associated input/output semantics.

In many cases, an input parameter specifies an effective address associated with a memory object inside or outside the EPC, the memory addressing semantics of these memory objects are also summarized in a separate table.

**EDECVIRTCHILD—Decrement VIRTCHILDCNT in SECS**

| Opcode/<br>Instruction            | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|-----------------------------------|-------|------------------------------|--------------------------|--|
| EAX = 00H<br>ENCLV[EDECVIRTCHILD] | IR    | V/V                          | EAX[5]                   | This leaf function decrements the SECS VIRTCHILDCNT field. |

**Instruction Operand Encoding**

| Op/En | EAX                | RBX                             | RCX                          |
|-------|--------------------|---------------------------------|------------------------------|
| IR    | EDECVIRTCHILD (In) | Address of an enclave page (In) | Address of an SECS page (In) |

**Description**

This instruction decrements the SECS VIRTCHILDCNT field. This instruction can only be executed when current privilege level is 0.

The content of RCX is an effective address of an EPC page. The DS segment is used to create linear address. Segment override is not supported.

**EDECVIRTCHILD Memory Parameter Semantics**

| EPCPAGE                                    | SECS                             |
|--|----------------------------------|
| Read/Write access permitted by Non Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

**EDECVIRTCHILD Faulting Conditions**

|  |   |
|--|---|
| A memory operand effective address is outside the DS segment limit (32b mode). | A page fault occurs in accessing memory operands.                             |
| DS segment is unusable (32b mode).   | RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).                 |
| A memory address is in a non-canonical form (64b mode).                        | RCX does not refer to an SECS page.   |
| A memory operand is not properly aligned.                                      | RBX does not refer to an enclave page associated with SECS referenced in RCX. |

**Concurrency Restrictions****Table 40-74. Base Concurrency Restrictions of EDECVIRTCHILD**

| Leaf          | Parameter       | Base Concurrency Restrictions |                       |                                    |
|---------------|-----------------|-------------------------------|-----------------------|------------------------------------|
|               |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| EDECVIRTCHILD | Target [DS:RBX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|               | SECS [DS:RCX]   | Concurrent                    |                       |                                    |

**Table 40-75. Additional Concurrency Restrictions of EDECVIRTCHILD**

| Leaf          | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|---------------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|               |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|               |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EDECVIRTCHILD | Target [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|               | SECS [DS:RCX]   | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EDECVIRTCHILD Operational Flow**

| Name             | Type             | Size (bits) | Description  |
|------------------|------------------|-------------|--|
| TMP_SECS         | Physical Address | 64          | Physical address of the SECS of the page being modified. |
| TMP_VIRTCHILDCNT | Integer          | 64          | Number of virtual child pages.                           |

**EDECVIRTCHILD Return Value in RAX**

| Error                 | Value | Description   |
|-----------------------|-------|---|
| No Error              | 0     | EDECVIRTCHILD Successful.                                       |
| SGX_EPC_PAGE_CONFLICT |       | Failure due to concurrent operation of another SGX instruction. |
| SGX_INVALID_COUNTER   |       | Attempt to decrement counter that is already zero.              |

(\* check alignment of DS:RBX \*)

```
IF (DS:RBX is not 4K aligned) THEN
    #GP(0); FI;
```

(\* check DS:RBX is an linear address of an EPC page \*)

```
IF (DS:RBX does not resolve within an EPC) THEN
    #PF(DS:RBX, PFEC.SGX); FI;
```

(\* check DS:RCX is an linear address of an EPC page \*)

```
IF (DS:RCX does not resolve within an EPC) THEN
    #PF(DS:RCX, PFEC.SGX); FI;
```

(\* Check the EPCPAGE for concurrency \*)

```
IF (EPCPAGE is being modified) THEN
    RFLAGS.ZF = 1;
    RAX = SGX_EPC_PAGE_CONFLICT;
    goto DONE;
FI;
```

(\* check that the EPC page is valid \*)

```
IF (EPCM(DS:RBX).VALID = 0) THEN
    #PF(DS:RBX, PFEC.SGX); FI;
```

(\* check that the EPC page has the correct type and that the back pointer matches the pointer passed as the pointer to parent \*)

```
IF ((EPCM(DS:RBX).PAGE_TYPE = PT_REG) or
    (EPCM(DS:RBX).PAGE_TYPE = PT_TCS) or
```

```

(EPCM(DS:RBX).PAGE_TYPE = PT_TRIM)
THEN
(* get the SECS of DS:RBX *)
TMP_SECS ← Address of SECS for (DS:RBX);
ELSE IF (EPCM(DS:RBX).PAGE_TYPE = PT_SECS) THEN
(* get the physical address of DS:RBX *)
TMP_SECS ← Physical_Address(DS:RBX);
ELSE
(* EDECVIRTCHILD called on page of incorrect type *)
#PF(DS:RBX, PFEC.SGX); FI;

IF (TMP_SECS ≠ Physical_Address(DS:RCX)) THEN
#GP(0); FI;

(* Atomically decrement virtchild counter and check for underflow *)
Locked_Decrement(SECS(TMP_SECS).VIRTCHILDCNT);
IF (There was an underflow) THEN
Locked_Increment(SECS(TMP_SECS).VIRTCHILDCNT);
RFLAGS.ZF ← 1;
RAX ? SGX_INVALID_COUNTER;
goto DONE;
FI;

RFLAGS.ZF ← 0;
RAX ← 0;

DONE:
(* clear flags *)
RFLAGS.CF ← 0;
RFLAGS.PF ← 0;
RFLAGS.AF ← 0;
RFLAGS.OF ← 0;
RFLAGS.SF ← 0;

```

### Flags Affected

ZF is set if EDECVIRTCHILD fails due to concurrent operation with another SGX instruction, or if there is a VIRTCHILDCNT underflow. Otherwise cleared.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <p>If a memory operand effective address is outside the DS segment limit.</p> <p>If DS segment is unusable.</p> <p>If a memory operand is not properly aligned.</p> <p>RBX does not refer to an enclave page associated with SECS referenced in RCX.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory operands.</p> <p>If RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).</p> <p>If RCX does not refer to an SECS page.</p>  |

**64-Bit Mode Exceptions**

- #GP(0)            If a memory address is in a non-canonical form.  
                  If a memory operand is not properly aligned.  
                  RBX does not refer to an enclave page associated with SECS referenced in RCX.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                  If RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).  
                  If RCX does not refer to an SECS page.

**EINCVIRTCHILD—Increment VIRTCHILDCNT in SECS**

| Opcode/<br>Instruction            | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description  |
|-----------------------------------|-------|------------------------------|--------------------------|--|
| EAX = 01H<br>ENCLV[EINCVIRTCHILD] | IR    | V/V                          | EAX[5]                   | This leaf function increments the SECS VIRTCHILDCNT field. |

**Instruction Operand Encoding**

| Op/En | EAX                | RBX                             | RCX                          |
|-------|--------------------|---------------------------------|------------------------------|
| IR    | EINCVIRTCHILD (In) | Address of an enclave page (In) | Address of an SECS page (In) |

**Description**

This instruction increments the SECS VIRTCHILDCNT field. This instruction can only be executed when the current privilege level is 0.

The content of RCX is an effective address of an EPC page. The DS segment is used to create a linear address. Segment override is not supported.

**EINCVIRTCHILD Memory Parameter Semantics**

| EPCPAGE                                    | SECS                             |
|--|----------------------------------|
| Read/Write access permitted by Non Enclave | Read access permitted by Enclave |

The instruction faults if any of the following:

**EINCVIRTCHILD Faulting Conditions**

|  |   |
|--|---|
| A memory operand effective address is outside the DS segment limit (32b mode). | A page fault occurs in accessing memory operands.                             |
| DS segment is unusable (32b mode).   | RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).                 |
| A memory address is in a non-canonical form (64b mode).                        | RCX does not refer to an SECS page.   |
| A memory operand is not properly aligned.                                      | RBX does not refer to an enclave page associated with SECS referenced in RCX. |

**Concurrency Restrictions****Table 40-76. Base Concurrency Restrictions of EINCVIRTCHILD**

| Leaf          | Parameter       | Base Concurrency Restrictions |                       |                                    |
|---------------|-----------------|-------------------------------|-----------------------|------------------------------------|
|               |                 | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| EINCVIRTCHILD | Target [DS:RBX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |
|               | SECS [DS:RCX]   | Concurrent                    |                       |                                    |

**Table 40-77. Additional Concurrency Restrictions of EINCVRTCHILD**

| Leaf         | Parameter       | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|--------------|-----------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|              |                 | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|              |                 | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| EINCVRTCHILD | Target [DS:RBX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |
|              | SECS [DS:RCX]   | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in EINCVRTCHILD Operational Flow**

| Name             | Type             | Size (bits) | Description  |
|------------------|------------------|-------------|--|
| TMP_SECS         | Physical Address | 64          | Physical address of the SECS of the page being modified. |
| TMP_VIRTCHILDCNT | Integer          | 64          | Number of virtual child pages.                           |

**EINCVRTCHILD Return Value in RAX**

| Error                 | Value | Description   |
|-----------------------|-------|---|
| No Error              | 0     | EINCVRTCHILD Successful.  |
| SGX_EPC_PAGE_CONFLICT |       | Failure due to concurrent operation of another SGX instruction. |
| SGX_INVALID_COUNTER   |       | Attempt to increment counter that will produce an overflow.     |

(\* check alignment of DS:RBX \*)

```
IF (DS:RBX is not 4K aligned) THEN
    #GP(0); FI;
```

(\* check DS:RBX is an linear address of an EPC page \*)

```
IF (DS:RBX does not resolve within an EPC) THEN
    #PF(DS:RBX, PFEC.SGX); FI;
```

(\* check DS:RCX is an linear address of an EPC page \*)

```
IF (DS:RCX does not resolve within an EPC) THEN
    #PF(DS:RCX, PFEC.SGX); FI;
```

(\* Check the EPCPAGE for concurrency \*)

```
IF (EPCPAGE is being modified) THEN
    RFLAGS.ZF = 1;
    RAX = SGX_EPC_PAGE_CONFLICT;
    goto DONE;
FI;
```

(\* check that the EPC page is valid \*)

```
IF (EPCM(DS:RBX).VALID = 0) THEN
    #PF(DS:RBX, PFEC.SGX); FI;
```

(\* check that the EPC page has the correct type and that the back pointer matches the pointer passed as the pointer to parent \*)

```
IF ((EPCM(DS:RBX).PAGE_TYPE = PT_REG) or
    (EPCM(DS:RBX).PAGE_TYPE = PT_TCS) or
```



```

(EPCM(DS:RBX).PAGE_TYPE = PT_TRIM)
THEN
  (* get the SECS of DS:RBX *)
  TMP_SECS ← Address of SECS for (DS:RBX);
ELSE IF (EPCM(DS:RBX).PAGE_TYPE = PT_SECS) THEN
  (* get the physical address of DS:RBX *)
  TMP_SECS ← Physical_Address(DS:RBX);
ELSE
  (* EINCVIRTCHILD called on page of incorrect type *)
  #PF(DS:RBX, PFEC.SGX); FI;

IF (TMP_SECS ≠ Physical_Address(DS:RCX)) THEN
  #GP(0); FI;

(* Atomically increment virtchild counter and check for overflow *)
Locked_Increment(SECS(TMP_SECS).VIRTCHILDCNT);
IF (There was an overflow) THEN
  Locked_Decrement(SECS(TMP_SECS).VIRTCHILDCNT);
  RFLAGS.ZF ← 1;
  RAX ← SGX_INVALID_COUNTER;
  goto DONE;
FI;

RFLAGS.ZF ← 0;
RAX ← 0;

DONE:
(* clear flags *)
RFLAGS.CF ← 0;
RFLAGS.PF ← 0;
RFLAGS.AF ← 0;
RFLAGS.OF ← 0;
RFLAGS.SF ← 0;

```

### Flags Affected

ZF is set if EINCVIRTCHILD fails due to concurrent operation with another SGX instruction, or if there is a VIRTCHILDCNT underflow; otherwise cleared.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | <p>If a memory operand effective address is outside the DS segment limit.</p> <p>If DS segment is unusable.</p> <p>If a memory operand is not properly aligned.</p> <p>RBX does not refer to an enclave page associated with SECS referenced in RCX.</p> |
| #PF(error code) | <p>If a page fault occurs in accessing memory operands.</p> <p>If RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).</p> <p>If RCX does not refer to an SECS page.</p>  |

**64-Bit Mode Exceptions**

- #GP(0)            If a memory address is in a non-canonical form.  
                  If a memory operand is not properly aligned.  
                  RBX does not refer to an enclave page associated with SECS referenced in RCX.
- #PF(error code)    If a page fault occurs in accessing memory operands.  
                  If RBX does not refer to an enclave page (REG, TCS, TRIM, SECS).  
                  If RCX does not refer to an SECS page.

## ESETCONTEXT—Set the ENCLAVECONTEXT Field in SECS

| Opcode/<br>Instruction          | Op/En | 64/32<br>bit Mode<br>Support | CPUID<br>Feature<br>Flag | Description   |
|---------------------------------|-------|------------------------------|--------------------------|---|
| EAX = 02H<br>ENCLV[ESETCONTEXT] | IR    | V/V                          | EAX[5]                   | This leaf function sets the ENCLAVECONTEXT field in SECS. |

### Instruction Operand Encoding

| Op/En | EAX              | RCX   | RDX                    |
|-------|------------------|---|------------------------|
| IR    | ESETCONTEXT (In) | Address of the destination EPC page<br>(In, EA) | Context Value (In, EA) |

### Description

The ESETCONTEXT leaf overwrites the ENCLAVECONTEXT field in the SECS. ECREATE and ELD of an SECS set the ENCLAVECONTEXT field in the SECS to the address of the SECS (for access later in ERDINFO). The ESETCONTEXT instruction allows a VMM to overwrite the default context value if necessary, for example, if the VMM is emulating ECREATE or ELD on behalf of the guest.

The content of RCX is an effective address of the SECS page to be updated, RDX contains the address pointing to the value to be stored in the SECS. The DS segment is used to create linear address. Segment override is not supported.

The instruction fails if:

- The operand is not properly aligned.
- RCX does not refer to an SECS page.

### ESETCONTEXT Memory Parameter Semantics

| EPCPAGE                          | CONTEXT                                    |
|----------------------------------|--|
| Read access permitted by Enclave | Read/Write access permitted by Non Enclave |

The instruction faults if any of the following:

### ESETCONTEXT Faulting Conditions

|  |   |
|--|---|
| A memory operand effective address is outside the DS segment limit (32b mode). | A memory operand is not properly aligned.         |
| DS segment is unusable (32b mode).   | A page fault occurs in accessing memory operands. |
| A memory address is in a non-canonical form (64b mode).                        |   |

### Concurrency Restrictions

**Table 40-78. Base Concurrency Restrictions of ESETCONTEXT**

| Leaf        | Parameter     | Base Concurrency Restrictions |                       |                                    |
|-------------|---------------|-------------------------------|-----------------------|------------------------------------|
|             |               | Access                        | On Conflict           | SGX_CONFLICT VM Exit Qualification |
| ESETCONTEXT | SECS [DS:RCX] | Shared                        | SGX_EPC_PAGE_CONFLICT |                                    |

**Table 40-79. Additional Concurrency Restrictions of ESETCONTEXT**

| Leaf        | Parameter     | Additional Concurrency Restrictions             |             |                          |             |                     |             |
|-------------|---------------|---|-------------|--------------------------|-------------|---------------------|-------------|
|             |               | vs. EACCEPT, EACCEPTCOPY, EMODPE, EMODPR, EMODT |             | vs. EADD, EEXTEND, EINIT |             | vs. ETRACK, ETRACKC |             |
|             |               | Access  | On Conflict | Access                   | On Conflict | Access              | On Conflict |
| ESETCONTEXT | SECS [DS:RCX] | Concurrent                                      |             | Concurrent               |             | Concurrent          |             |

**Operation**

**Temp Variables in ESETCONTEXT Operational Flow**

| Name        | Type             | Size (bits) | Description  |
|-------------|------------------|-------------|--|
| TMP_SECS    | Physical Address | 64          | Physical address of the SECS of the page being modified. |
| TMP_CONTEXT | CONTEXT          | 64          | Data Value of CONTEXT.                                   |

**ESETCONTEXT Return Value in RAX**

| Error                 | Value | Description   |
|-----------------------|-------|---|
| No Error              | 0     | ESETCONTEXT Successful.   |
| SGX_EPC_PAGE_CONFLICT |       | Failure due to concurrent operation of another SGX instruction. |

(\* check alignment of the EPCPAGE (RCX) \*)

```
IF (DS:RCX is not 4KByte Aligned) THEN
    #GP(0); FI;
```

(\* check that EPCPAGE (DS:RCX) is the address of an EPC page \*)

```
IF (DS:RCX does not resolve within an EPC) THEN
    #PF(DS:RCX, PFEC.SGX); FI;
```

(\* check alignment of the CONTEXT field (RDX) \*)

```
IF (DS:RDX is not 8Byte Aligned) THEN
    #GP(0); FI;
```

(\* Load CONTEXT into local variable \*)

```
TMP_CONTEXT ← DS:RDX
```

(\* Check the EPC page for concurrency \*)

```
IF (EPC page is being modified) THEN
    RFLAGS.ZF ← 1;
    RFLAGS.CF ← 0;
    RAX ← SGX_EPC_PAGE_CONFLICT;
    goto DONE;
FI;
```

(\* check page validity \*)

```
IF (EPCM(DS:RCX).VALID = 0) THEN
    #PF(DS:RCX, PFEC.SGX);
    goto DONE;
FI;
```

```
(* check EPC page is an SECS page *)
IF (EPCM(DS:RCX).PT is not PT_SECS) THEN
  #PF(DS:RCX, PFEC.SGX);
  goto DONE;
FI;
```

```
(* load the context value into SECS(DS:RCX).ENCLAVECONTEXT *)
SECS(DS:RCX).ENCLAVECONTEXT ← TMP_CONTEXT;
```

```
RAX ← 0;
RFLAGS.ZF ← 0;
```

```
DONE:
(* clear flags *)
RFLAGS.CF,PF,AF,OF,SF ← 0;
```

### Flags Affected

ZF is set if ESETCONTEXT fails due to concurrent operation with another SGX instruction; otherwise cleared. CF, PF, AF, OF and SF are cleared.

### Protected Mode Exceptions

|                 |  |
|-----------------|--|
| #GP(0)          | If a memory operand effective address is outside the DS segment limit.<br>If DS segment is unusable.<br>If a memory operand is not properly aligned. |
| #PF(error code) | If a page fault occurs in accessing memory operands.   |

### 64-Bit Mode Exceptions

|                 |   |
|-----------------|---|
| #GP(0)          | If a memory address is in a non-canonical form.<br>If a memory operand is not properly aligned. |
| #PF(error code) | If a page fault occurs in accessing memory operands.  |



## 16. Updates to Chapter 1, Volume 4

Change bars show changes to Chapter 1 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers*.

-----  
Changes to this chapter: Updates to processors covered by manual; added 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series.

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers* (order number 335592) is part of a set that describes the architecture and programming environment of Intel® 64 and IA-32 architecture processors. Other volumes in this set are:

- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture* (order number 253665).
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D: Instruction Set Reference* (order numbers 253666, 253667, 326018 and 334569).
- *The Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D: System Programming Guide* (order numbers 253668, 253669, 326019 and 332831).

The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, describes the basic architecture and programming environment of Intel 64 and IA-32 processors. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 2A, 2B, 2C & 2D*, describe the instruction set of the processor and the opcode structure. These volumes apply to application programmers and to programmers who write operating systems or executives. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A, 3B, 3C & 3D*, describe the operating-system support environment of Intel 64 and IA-32 processors. These volumes target operating-system and BIOS designers. In addition, *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B*, and *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C* address the programming environment for classes of software that host operating systems. The *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4*, describes the model-specific registers of Intel 64 and IA-32 processors.

## 1.1 INTEL® 64 AND IA-32 PROCESSORS COVERED IN THIS MANUAL

This manual set includes information pertaining primarily to the most recent Intel 64 and IA-32 processors, which include:

- Pentium® processors
- P6 family processors
- Pentium® 4 processors
- Pentium® M processors
- Intel® Xeon® processors
- Pentium® D processors
- Pentium® processor Extreme Editions
- 64-bit Intel® Xeon® processors
- Intel® Core™ Duo processor
- Intel® Core™ Solo processor
- Dual-Core Intel® Xeon® processor LV
- Intel® Core™2 Duo processor
- Intel® Core™2 Quad processor Q6000 series
- Intel® Xeon® processor 3000, 3200 series
- Intel® Xeon® processor 5000 series
- Intel® Xeon® processor 5100, 5300 series
- Intel® Core™2 Extreme processor X7000 and X6800 series
- Intel® Core™2 Extreme QX6000 series
- Intel® Xeon® processor 7100 series



## ABOUT THIS MANUAL

- Intel® Pentium® Dual-Core processor
- Intel® Xeon® processor 7200, 7300 series
- Intel® Core™2 Extreme QX9000 series
- Intel® Xeon® processor 5200, 5400, 7400 series
- Intel® Core™2 Extreme processor QX9000 and X9000 series
- Intel® Core™2 Quad processor Q9000 series
- Intel® Core™2 Duo processor E8000, T9000 series
- Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are built from 45 nm and 32 nm processes.
- Intel® Core™ i7 processor
- Intel® Core™ i5 processor
- Intel® Xeon® processor E7-8800/4800/2800 product families
- Intel® Core™ i7-3930K processor
- 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series
- Intel® Xeon® processor E3-1200 product family
- Intel® Xeon® processor E5-2400/1400 product family
- Intel® Xeon® processor E5-4600/2600/1600 product family
- 3rd generation Intel® Core™ processors
- Intel® Xeon® processor E3-1200 v2 product family
- Intel® Xeon® processor E5-2400/1400 v2 product families
- Intel® Xeon® processor E5-4600/2600/1600 v2 product families
- Intel® Xeon® processor E7-8800/4800/2800 v2 product families
- 4th generation Intel® Core™ processors
- The Intel® Core™ M processor family
- Intel® Core™ i7-59xx Processor Extreme Edition
- Intel® Core™ i7-49xx Processor Extreme Edition
- Intel® Xeon® processor E3-1200 v3 product family
- Intel® Xeon® processor E5-2600/1600 v3 product families
- 5th generation Intel® Core™ processors
- Intel® Xeon® processor D-1500 product family
- Intel® Xeon® processor E5 v4 family
- Intel® Atom™ processor X7-Z8000 and X5-Z8000 series
- Intel® Atom™ processor Z3400 series
- Intel® Atom™ processor Z3500 series
- 6th generation Intel® Core™ processors
- Intel® Xeon® processor E3-1500m v5 product family
- 7th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series
- Intel® Xeon® Processor Scalable Family
- 8th generation Intel® Core™ processors
- Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series

P6 family processors are IA-32 processors based on the P6 family microarchitecture. This includes the Pentium® Pro, Pentium® II, Pentium® III, and Pentium® III Xeon® processors.

The Pentium® 4, Pentium® D, and Pentium® processor Extreme Editions are based on the Intel NetBurst® microarchitecture. Most early Intel® Xeon® processors are based on the Intel NetBurst® microarchitecture. Intel Xeon processor 5000, 7100 series are based on the Intel NetBurst® microarchitecture.

The Intel® Core™ Duo, Intel® Core™ Solo and dual-core Intel® Xeon® processor LV are based on an improved Pentium® M processor microarchitecture.

The Intel® Xeon® processor 3000, 3200, 5100, 5300, 7200, and 7300 series, Intel® Pentium® dual-core, Intel® Core™2 Duo, Intel® Core™2 Quad, and Intel® Core™2 Extreme processors are based on Intel® Core™ microarchitecture.

The Intel® Xeon® processor 5200, 5400, 7400 series, Intel® Core™2 Quad processor Q9000 series, and Intel® Core™2 Extreme processors QX9000, X9000 series, Intel® Core™2 processor E8000 series are based on Enhanced Intel® Core™ microarchitecture.

The Intel® Atom™ processors 200, 300, D400, D500, D2000, N200, N400, N2000, E2000, Z500, Z600, Z2000, C1000 series are based on the Intel® Atom™ microarchitecture and supports Intel 64 architecture.

P6 family, Pentium® M, Intel® Core™ Solo, Intel® Core™ Duo processors, dual-core Intel® Xeon® processor LV, and early generations of Pentium 4 and Intel Xeon processors support IA-32 architecture. The Intel® Atom™ processor Z5xx series support IA-32 architecture.

The Intel® Xeon® processor 3000, 3200, 5000, 5100, 5200, 5300, 5400, 7100, 7200, 7300, 7400 series, Intel® Core™2 Duo, Intel® Core™2 Extreme, Intel® Core™2 Quad processors, Pentium® D processors, Pentium® Dual-Core processor, newer generations of Pentium 4 and Intel Xeon processor family support Intel® 64 architecture.

The Intel® Core™ i7 processor and Intel® Xeon® processor 3400, 5500, 7500 series are based on 45 nm Intel® microarchitecture code name Nehalem. Intel® microarchitecture code name Westmere is a 32 nm version of Intel® microarchitecture code name Nehalem. Intel® Xeon® processor 5600 series, Intel Xeon processor E7 and various Intel Core i7, i5, i3 processors are based on Intel® microarchitecture code name Westmere. These processors support Intel 64 architecture.

The Intel® Xeon® processor E5 family, Intel® Xeon® processor E3-1200 family, Intel® Xeon® processor E7-8800/4800/2800 product families, Intel® Core™ i7-3930K processor, and 2nd generation Intel® Core™ i7-2xxx, Intel® Core™ i5-2xxx, Intel® Core™ i3-2xxx processor series are based on the Intel® microarchitecture code name Sandy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E7-8800/4800/2800 v2 product families, Intel® Xeon® processor E3-1200 v2 product family and 3rd generation Intel® Core™ processors are based on the Intel® microarchitecture code name Ivy Bridge and support Intel 64 architecture.

The Intel® Xeon® processor E5-4600/2600/1600 v2 product families, Intel® Xeon® processor E5-2400/1400 v2 product families and Intel® Core™ i7-49xx Processor Extreme Edition are based on the Intel® microarchitecture code name Ivy Bridge-E and support Intel 64 architecture.

The Intel® Xeon® processor E3-1200 v3 product family and 4th Generation Intel® Core™ processors are based on the Intel® microarchitecture code name Haswell and support Intel 64 architecture.

The Intel® Xeon® processor E5-2600/1600 v3 product families and the Intel® Core™ i7-59xx Processor Extreme Edition are based on the Intel® microarchitecture code name Haswell-E and support Intel 64 architecture.

The Intel® Atom™ processor Z8000 series is based on the Intel microarchitecture code name Airmont.

The Intel® Atom™ processor Z3400 series and the Intel® Atom™ processor Z3500 series are based on the Intel microarchitecture code name Silvermont.

The Intel® Core™ M processor family, 5th generation Intel® Core™ processors, Intel® Xeon® processor D-1500 product family and the Intel® Xeon® processor E5 v4 family are based on the Intel® microarchitecture code name Broadwell and support Intel 64 architecture.

The Intel® Xeon® Processor Scalable Family, Intel® Xeon® processor E3-1500m v5 product family and 6th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Skylake and support Intel 64 architecture.

The 7th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Kaby Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series is based on the Intel® microarchitecture code name Knights Landing and supports Intel 64 architecture.

The 8th generation Intel® Core™ processors are based on the Intel® microarchitecture code name Coffee Lake and support Intel 64 architecture.

The Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series is based on the Intel® microarchitecture code name Knights Mill and supports Intel 64 architecture.

IA-32 architecture is the instruction set architecture and programming environment for Intel's 32-bit microprocessors. Intel® 64 architecture is the instruction set architecture and programming environment which is the superset of Intel's 32-bit and 64-bit architectures. It is compatible with the IA-32 architecture.

## 1.2 OVERVIEW OF THE SYSTEM PROGRAMMING GUIDE

A description of this manual's content follows:

**Chapter 1 — About This Manual.** Gives an overview of all eight volumes of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*. It also describes the notational conventions in these manuals and lists related Intel manuals and documentation of interest to programmers and hardware designers.

**Chapter 2 — Model-Specific Registers (MSRs).** Lists the MSRs available in the Pentium processors, the P6 family processors, the Pentium 4, Intel Xeon, Intel Core Solo, Intel Core Duo processors, Intel Core 2 processor family, and Intel Atom processors, and describes their functions.

## 1.3 NOTATIONAL CONVENTIONS

This manual uses specific notation for data-structure formats, for symbolic representation of instructions, and for hexadecimal and binary numbers. A review of this notation makes the manual easier to read.

### 1.3.1 Bit and Byte Order

In illustrations of data structures in memory, smaller addresses appear toward the bottom of the figure; addresses increase toward the top. Bit positions are numbered from right to left. The numerical value of a set bit is equal to two raised to the power of the bit position. Intel 64 and IA-32 processors are "little endian" machines; this means the bytes of a word are numbered starting from the least significant byte. Figure 1-1 illustrates these conventions.

### 1.3.2 Reserved Bits and Software Compatibility

In many register and memory layout descriptions, certain bits are marked as reserved. When bits are marked as reserved, it is essential for compatibility with future processors that software treat these bits as having a future, though unknown, effect. The behavior of reserved bits should be regarded as not only undefined, but unpredictable. Software should follow these guidelines in dealing with reserved bits:

- Do not depend on the states of any reserved bits when testing the values of registers which contain such bits. Mask out the reserved bits before testing.
- Do not depend on the states of any reserved bits when storing to memory or to a register.
- Do not depend on the ability to retain information written into any reserved bits.
- When loading a register, always load the reserved bits with the values indicated in the documentation, if any, or reload them with values previously read from the same register.

#### NOTE

Avoid any software dependence upon the state of reserved bits in Intel 64 and IA-32 registers. Depending upon the values of reserved register bits will make software dependent upon the unspecified manner in which the processor handles these bits. Programs that depend upon reserved values risk incompatibility with future processors.



The processor also supports segmented addressing. This is a form of addressing where a program may have many independent address spaces, called **segments**. For example, a program can keep its code (instructions) and stack in separate segments. Code addresses would always refer to the code space, and stack addresses would always refer to the stack space. The following notation is used to specify a byte address within a segment:

Segment-register:Byte-address

For example, the following segment address identifies the byte at address FF79H in the segment pointed to by the DS register:

DS:FF79H

The following segment address identifies an instruction address in the code segment. The CS register points to the code segment and the EIP register contains the address of the instruction.

CS:EIP

### 1.3.6 Syntax for CPUID, CR, and MSR Values

Obtain feature flags, status, and system information by using the CPUID instruction, by checking control register bits, and by reading model-specific registers. We are moving toward a single syntax to represent this type of information. See Figure 1-2.




Figure 1-2. Syntax for CPUID, CR, and MSR Data Presentation

### 1.3.7 Exceptions

An exception is an event that typically occurs when an instruction causes an error. For example, an attempt to divide by zero generates an exception. However, some exceptions, such as breakpoints, occur under other conditions. Some types of exceptions may provide error codes. An error code reports additional information about the error. An example of the notation used to show an exception and error code is shown below:

```
#PF(fault code)
```

This example refers to a page-fault exception under conditions where an error code naming a type of fault is reported. Under some conditions, exceptions which produce error codes may not be able to report an accurate code. In this case, the error code is zero, as shown below for a general-protection exception:

```
#GP(0)
```

## 1.4 RELATED LITERATURE

Literature related to Intel 64 and IA-32 processors is listed and viewable on-line at:

<http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>

See also:

- The data sheet for a particular Intel 64 or IA-32 processor
- The specification update for a particular Intel 64 or IA-32 processor
- Intel® C++ Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Fortran Compiler documentation and online help:  
<http://software.intel.com/en-us/articles/intel-compilers/>
- Intel® Software Development Tools:  
<https://software.intel.com/en-us/intel-sdp-home>
- Intel® 64 and IA-32 Architectures Software Developer's Manual (in one, four or ten volumes):  
<https://software.intel.com/en-us/articles/intel-sdm>
- Intel® 64 and IA-32 Architectures Optimization Reference Manual:  
<https://software.intel.com/en-us/articles/intel-sdm#optimization>
- Intel 64 Architecture x2APIC Specification:  
<http://www.intel.com/content/www/us/en/architecture-and-technology/64-architecture-x2apic-specification.html>
- Intel® Trusted Execution Technology Measured Launched Environment Programming Guide:  
<http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>
- Developing Multi-threaded Applications: A Platform Consistent Approach:  
<https://software.intel.com/sites/default/files/article/147714/51534-developing-multithreaded-applications.pdf>
- Using Spin-Loops on Intel® Pentium® 4 Processor and Intel® Xeon® Processor:  
<https://software.intel.com/sites/default/files/22/30/25602>
- Performance Monitoring Unit Sharing Guide  
<http://software.intel.com/file/30388>

Literature related to selected features in future Intel processors are available at:

- Intel® Architecture Instruction Set Extensions Programming Reference  
<https://software.intel.com/en-us/isa-extensions>
- Intel® Software Guard Extensions (Intel® SGX) Programming Reference  
<https://software.intel.com/en-us/isa-extensions/intel-sgx>

## ABOUT THIS MANUAL

More relevant links are:

- Intel® Developer Zone:  
<https://software.intel.com/en-us>
- Developer centers:  
<http://www.intel.com/content/www/us/en/hardware-developers/developer-centers.html>
- Processor support general link:  
<http://www.intel.com/support/processors/>
- Intel® Hyper-Threading Technology (Intel® HT Technology):  
<http://www.intel.com/technology/platform-technology/hyper-threading/index.htm>

## 17. Updates to Chapter 2, Volume 4

Change bars show changes to Chapter 2 of the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 4: Model-Specific Registers*.

-----  
Changes to this chapter: Updates to add 8th generation Intel® Core™ processors and Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series to existing tables. Many minor typo corrections (these are not marked by change bars as they are not technical changes and numerous in number).



## CHAPTER 2

# MODEL-SPECIFIC REGISTERS (MSRS)

This chapter lists MSRs across Intel processor families. All MSRs listed can be read with the RDMSR and written with the WRMSR instructions.

Register addresses are given in both hexadecimal and decimal. The register name is the mnemonic register name and the bit description describes individual bits in registers.

Model specific registers and its bit-fields may be supported for a finite range of processor families/models. To distinguish between different processor family and/or models, software must use CPUID.01H leaf function to query the combination of DisplayFamily and DisplayModel to determine model-specific availability of MSRs (see CPUID instruction in Chapter 3, "Instruction Set Reference, A-L" in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A*). Table 2-1 lists the signature values of DisplayFamily and DisplayModel for various processor families or processor number series.

**Table 2-1. CPUID Signature Values of DisplayFamily\_DisplayModel**

| DisplayFamily_DisplayModel | Processor Families/Processor Number Series   |
|----------------------------|--|
| 06_85H                     | Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series based on Knights Mill microarchitecture   |
| 06_57H                     | Intel® Xeon Phi™ Processor 3200, 5200, 7200 Series based on Knights Landing microarchitecture  |
| 06_66H                     | Future Intel® Core™ processors based on Cannon Lake microarchitecture  |
| 06_8EH, 06_9EH             | 7th generation Intel® Core™ processors based on Kaby Lake microarchitecture and 8th generation Intel® Core™ processors based on Coffee Lake microarchitecture                                  |
| 06_55H                     | Intel® Xeon® Processor Scalable Family based on Skylake microarchitecture  |
| 06_4EH, 06_5EH             | 6th generation Intel Core processors and Intel Xeon processor E3-1500m v5 product family and E3-1200 v5 product family based on Skylake microarchitecture                                      |
| 06_56H                     | Intel Xeon processor D-1500 product family based on Broadwell microarchitecture  |
| 06_4FH                     | Intel Xeon processor E5 v4 Family based on Broadwell microarchitecture, Intel Xeon processor E7 v4 Family, Intel Core i7-69xx Processor Extreme Edition  |
| 06_47H                     | 5th generation Intel Core processors, Intel Xeon processor E3-1200 v4 product family based on Broadwell microarchitecture  |
| 06_3DH                     | Intel Core M-5xxx Processor, 5th generation Intel Core processors based on Broadwell microarchitecture   |
| 06_3FH                     | Intel Xeon processor E5-4600/2600/1600 v3 product families, Intel Xeon processor E7 v3 product families based on Haswell-E microarchitecture, Intel Core i7-59xx Processor Extreme Edition     |
| 06_3CH, 06_45H, 06_46H     | 4th Generation Intel Core processor and Intel Xeon processor E3-1200 v3 product family based on Haswell microarchitecture  |
| 06_3EH                     | Intel Xeon processor E7-8800/4800/2800 v2 product families based on Ivy Bridge-E microarchitecture   |
| 06_3EH                     | Intel Xeon processor E5-2600/1600 v2 product families and Intel Xeon processor E5-2400 v2 product family based on Ivy Bridge-E microarchitecture, Intel Core i7-49xx Processor Extreme Edition |
| 06_3AH                     | 3rd Generation Intel Core Processor and Intel Xeon processor E3-1200 v2 product family based on Ivy Bridge microarchitecture   |
| 06_2DH                     | Intel Xeon processor E5 Family based on Intel microarchitecture code name Sandy Bridge, Intel Core i7-39xx Processor Extreme Edition   |
| 06_2FH                     | Intel Xeon Processor E7 Family   |
| 06_2AH                     | Intel Xeon processor E3-1200 product family; 2nd Generation Intel Core i7, i5, i3 Processors 2xxx Series   |
| 06_2EH                     | Intel Xeon processor 7500, 6500 series   |

**Table 2-1. CPUID Signature (Contd.)Values of DisplayFamily\_DisplayModel (Contd.)**

| DisplayFamily_DisplayModel             | Processor Families/Processor Number Series   |
|--|--|
| 06_25H, 06_2CH                         | Intel Xeon processors 3600, 5600 series, Intel Core i7, i5 and i3 Processors   |
| 06_1EH, 06_1FH                         | Intel Core i7 and i5 Processors  |
| 06_1AH                                 | Intel Core i7 Processor, Intel Xeon processor 3400, 3500, 5500 series  |
| 06_1DH                                 | Intel Xeon processor MP 7400 series  |
| 06_17H                                 | Intel Xeon processor 3100, 3300, 5200, 5400 series, Intel Core 2 Quad processors 8000, 9000 series   |
| 06_0FH                                 | Intel Xeon processor 3000, 3200, 5100, 5300, 7300 series, Intel Core 2 Quad processor 6000 series, Intel Core 2 Extreme 6000 series, Intel Core 2 Duo 4000, 5000, 6000, 7000 series processors, Intel Pentium dual-core processors |
| 06_0EH                                 | Intel Core Duo, Intel Core Solo processors   |
| 06_0DH                                 | Intel Pentium M processor  |
| 06_7AH                                 | Intel® Atom™ processors based on Goldmont Plus Microarchitecture   |
| 06_5FH                                 | Intel Atom processors based on Goldmont Microarchitecture (code name Denverton)  |
| 06_5CH                                 | Intel Atom processors based on Goldmont Microarchitecture  |
| 06_4CH                                 | Intel Atom processor X7-Z8000 and X5-Z8000 series based on Airmont Microarchitecture   |
| 06_5DH                                 | Intel Atom processor X3-C3000 based on Silvermont Microarchitecture  |
| 06_5AH                                 | Intel Atom processor Z3500 series  |
| 06_4AH                                 | Intel Atom processor Z3400 series  |
| 06_37H                                 | Intel Atom processor E3000 series, Z3600 series, Z3700 series  |
| 06_4DH                                 | Intel Atom processor C2000 series  |
| 06_36H                                 | Intel Atom processor S1000 Series  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H | Intel Atom processor family, Intel Atom processor D2000, N2000, E2000, Z2000, C1000 series   |
| 0F_06H                                 | Intel Xeon processor 7100, 5000 Series, Intel Xeon Processor MP, Intel Pentium 4, Pentium D processors   |
| 0F_03H, 0F_04H                         | Intel Xeon processor, Intel Xeon processor MP, Intel Pentium 4, Pentium D processors   |
| 06_09H                                 | Intel Pentium M processor  |
| 0F_02H                                 | Intel Xeon Processor, Intel Xeon processor MP, Intel Pentium 4 processors  |
| 0F_0H, 0F_01H                          | Intel Xeon Processor, Intel Xeon processor MP, Intel Pentium 4 processors  |
| 06_7H, 06_08H, 06_0AH, 06_0BH          | Intel Pentium III Xeon processor, Intel Pentium III processor  |
| 06_03H, 06_05H                         | Intel Pentium II Xeon processor, Intel Pentium II processor  |
| 06_01H                                 | Intel Pentium Pro processor  |
| 05_01H, 05_02H, 05_04H                 | Intel Pentium processor, Intel Pentium processor with MMX Technology   |

The Intel® Quark™ SoC X1000 processor can be identified by the signature of DisplayFamily\_DisplayModel = 05\_09H and SteppingID = 0

## 2.1 ARCHITECTURAL MSRS

Many MSRs have carried over from one generation of IA-32 processors to the next and to Intel 64 processors. A subset of MSRs and associated bit fields, which do not change on future processor generations, are now considered architectural MSRs. For historical reasons (beginning with the Pentium 4 processor), these “architectural MSRs” were given the prefix “IA32\_”. Table 2-2 lists the architectural MSRs, their addresses, their current names, their names in previous IA-32 processors, and bit fields that are considered architectural. MSR addresses outside Table 2-2 and certain bit fields in an MSR address that may overlap with architectural MSR addresses are model-specific.

Code that accesses a machine specified MSR and that is executed on a processor that does not support that MSR will generate an exception.

Architectural MSR or individual bit fields in an architectural MSR may be introduced or transitioned at the granularity of certain processor family/model or the presence of certain CPUID feature flags. The right-most column of Table 2-2 provides information on the introduction of each architectural MSR or its individual fields. This information is expressed either as signature values of "DF\_DM" (see Table 2-1) or via CPUID flags.

Certain bit field position may be related to the maximum physical address width, the value of which is expressed as "MAXPHYADDR" in Table 2-2. "MAXPHYADDR" is reported by CPUID.8000\_0008H leaf.

MSR address range between 40000000H - 400000FFH is marked as a specially reserved range. All existing and future processors will not implement any features using any MSR in this range.

**Table 2-2. IA-32 Architectural MSRs**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                    |
|------------------|---------|--|--|----------------------------|
| Hex              | Decimal |  |  |                            |
| 0H               | 0       | IA32_P5_MC_ADDR (P5_MC_ADDR)                             | See Section 2.22, "MSRs in Pentium Processors."  | Pentium Processor (05_01H) |
| 1H               | 1       | IA32_P5_MC_TYPE (P5_MC_TYPE)                             | See Section 2.22, "MSRs in Pentium Processors."  | DF_DM = 05_01H             |
| 6H               | 6       | IA32_MONITOR_FILTER_SIZE                                 | See Section 8.10.5, "Monitor/Mwait Address Range Determination."   | 0F_03H                     |
| 10H              | 16      | IA32_TIME_STAMP_COUNTER<br>(TSC)                         | See Section 17.17, "Time-Stamp Counter."   | 05_01H                     |
| 17H              | 23      | IA32_PLATFORM_ID<br>(MSR_PLATFORM_ID)                    | Platform ID (RO)<br>The operating system can use this MSR to determine "slot" information for the processor and the proper microcode update to load.   | 06_01H                     |
|                  |         | 49:0   | Reserved   |                            |
|                  |         | 52:50  | Platform Id (RO)<br>Contains information concerning the intended platform for the processor.<br>52 51 50<br>0 0 0 Processor Flag 0<br>0 0 1 Processor Flag 1<br>0 1 0 Processor Flag 2<br>0 1 1 Processor Flag 3<br>1 0 0 Processor Flag 4<br>1 0 1 Processor Flag 5<br>1 1 0 Processor Flag 6<br>1 1 1 Processor Flag 7 |                            |
|                  |         | 63:53  | Reserved   |                            |
| 1BH              | 27      | IA32_APIC_BASE<br>(APIC_BASE)                            | This register holds the APIC base address, permitting the relocation of the APIC memory map. See Section 10.4.4, "Local APIC Status and Location" and Section 10.4.5, "Relocating the Local APIC Registers".   | 06_01H                     |
|                  |         | 7:0  | Reserved   |                            |
|                  |         | 8  | BSP flag (R/W)   |                            |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment   |
|------------------|---------|--|---|---|
| Hex              | Decimal |  |   |   |
|                  |         | 9  | Reserved  |   |
|                  |         | 10   | Enable x2APIC mode.   | 06_1AH  |
|                  |         | 11   | APIC Global Enable (R/W)  |   |
|                  |         | (MAXPHYADDR - 1):12                                      | APIC Base (R/W)   |   |
|                  |         | 63: MAXPHYADDR   | Reserved  |   |
| 3AH              | 58      | IA32_FEATURE_CONTROL                                     | Control Features in Intel 64 Processor (R/W)  | If any one enumeration condition for defined bit field holds.                             |
|                  |         | 0  | Lock bit (R/WO): (1 = locked). When set, locks this MSR from being written; writes to this bit will result in GP(0).<br><br>Note: Once the Lock bit is set, the contents of this register cannot be modified. Therefore the lock bit must be set after configuring support for Intel Virtualization Technology and prior to transferring control to an option ROM or the OS. Hence, once the Lock bit is set, the entire IA32_FEATURE_CONTROL contents are preserved across RESET when PWRGOOD is not deasserted. | If any one enumeration condition for defined bit field position greater than bit 0 holds. |
|                  |         | 1  | Enable VMX inside SMX operation (R/WL): This bit enables a system executive to use VMX in conjunction with SMX to support Intel® Trusted Execution Technology.<br><br>BIOS must set this bit only when the CPUID function 1 returns VMX feature flag and SMX feature flag set (ECX bits 5 and 6 respectively).  | If CPUID.01H:ECX[5] = 1 && CPUID.01H:ECX[6] = 1   |
|                  |         | 2  | Enable VMX outside SMX operation (R/WL): This bit enables VMX for a system executive that does not require SMX.<br><br>BIOS must set this bit only when the CPUID function 1 returns the VMX feature flag set (ECX bit 5).  | If CPUID.01H:ECX[5] = 1   |
|                  |         | 7:3  | Reserved  |   |
|                  |         | 14:8   | SENTER Local Function Enables (R/WL): When set, each bit in the field represents an enable control for a corresponding SENTER function. This bit is supported only if CPUID.1:ECX.[bit 6] is set.   | If CPUID.01H:ECX[6] = 1   |
|                  |         | 15   | SENTER Global Enable (R/WL): This bit must be set to enable SENTER leaf functions. This bit is supported only if CPUID.1:ECX.[bit 6] is set.  | If CPUID.01H:ECX[6] = 1   |
|                  |         | 16   | Reserved  |   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                                 |
|------------------|---------|--|---|---|
| Hex              | Decimal |  |   |   |
|                  |         | 17   | SGX Launch Control Enable (R/WL): This bit must be set to enable runtime reconfiguration of SGX Launch Control via the IA32_SGXLEPUBKEYHASHn MSR.   | If CPUID.(EAX=07H, ECX=0H): ECX[30] = 1 |
|                  |         | 18   | SGX Global Enable (R/WL): This bit must be set to enable SGX leaf functions.  | If CPUID.(EAX=07H, ECX=0H): EBX[2] = 1  |
|                  |         | 19   | Reserved  |   |
|                  |         | 20   | LMCE On (R/WL): When set, system software can program the MSRs associated with LMCE to configure delivery of some machine check exceptions to a single logical processor.   | If IA32_MCG_CAP[27] = 1                 |
|                  |         | 63:21  | Reserved  |   |
| 3BH              | 59      | IA32_TSC_ADJUST  | Per Logical Processor TSC Adjust (R/Write to clear)   | If CPUID.(EAX=07H, ECX=0H): EBX[1] = 1  |
|                  |         | 63:0   | THREAD_ADJUST:<br>Local offset value of the IA32_TSC for a logical processor. Reset value is zero. A write to IA32_TSC will modify the local offset in IA32_TSC_ADJUST and the content of IA32_TSC, but does not affect the internal invariant TSC hardware.  |   |
| 79H              | 121     | IA32_BIOS_UPDT_TRIG<br>(BIOS_UPDT_TRIG)                  | BIOS Update Trigger (W)<br>Executing a WRMSR instruction to this MSR causes a microcode update to be loaded into the processor. See Section 9.11.6, "Microcode Update Loader."<br>A processor may prevent writing to this MSR when loading guest states on VM entries or saving guest states on VM exits. | 06_01H                                  |
| 8BH              | 139     | IA32_BIOS_SIGN_ID<br>(BIOS_SIGN/BBL_CR_D3)               | BIOS Update Signature (RO)<br>Returns the microcode update signature following the execution of CPUID.01H.<br>A processor may prevent writing to this MSR when loading guest states on VM entries or saving guest states on VM exits.   | 06_01H                                  |
|                  |         | 31:0   | Reserved  |   |
|                  |         | 63:32  | It is recommended that this field be pre-loaded with zero prior to executing CPUID. If the field remains zero following the execution of CPUID, this indicates that no microcode update is loaded. Any non-zero value is the microcode update signature.  |   |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment   |
|------------------|---------|--|---|---|
| Hex              | Decimal |  |   |   |
| 8CH              | 140     | IA32_SGXLEPUBKEYHASH0                                    | IA32_SGXLEPUBKEYHASH[63:0] (R/W)<br>Bits 63:0 of the SHA256 digest of the SIGSTRUCT.MODULUS for SGX Launch Enclave. On reset, the default value is the digest of Intel's signing key.       | Read permitted If<br>CPUID.(EAX=12H,ECX=0H):<br>EAX[0]=1 &&<br>CPUID.(EAX=07H,<br>ECX=0H):ECX[30]=1.<br><br>Write permitted if<br>CPUID.(EAX=12H,ECX=0H):<br>EAX[0]=1 &&<br>IA32_FEATURE_CONTROL[<br>17] = 1 &&<br>IA32_FEATURE_CONTROL[<br>0] = 1. |
| 8DH              | 141     | IA32_SGXLEPUBKEYHASH1                                    | IA32_SGXLEPUBKEYHASH[127:64] (R/W)<br>Bits 127:64 of the SHA256 digest of the SIGSTRUCT.MODULUS for SGX Launch Enclave. On reset, the default value is the digest of Intel's signing key.   |   |
| 8EH              | 142     | IA32_SGXLEPUBKEYHASH2                                    | IA32_SGXLEPUBKEYHASH[191:128] (R/W)<br>Bits 191:128 of the SHA256 digest of the SIGSTRUCT.MODULUS for SGX Launch Enclave. On reset, the default value is the digest of Intel's signing key. |   |
| 8FH              | 143     | IA32_SGXLEPUBKEYHASH3                                    | IA32_SGXLEPUBKEYHASH[255:192] (R/W)<br>Bits 255:192 of the SHA256 digest of the SIGSTRUCT.MODULUS for SGX Launch Enclave. On reset, the default value is the digest of Intel's signing key. |   |
| 9BH              | 155     | IA32_SMM_MONITOR_CTL                                     | SMM Monitor Configuration (R/W)   | If CPUID.01H: ECX[5]=1   <br>CPUID.01H: ECX[6] = 1  |
|                  |         | 0  | Valid (R/W)   |   |
|                  |         | 1  | Reserved  |   |
|                  |         | 2  | Controls SMI unblocking by VMXOFF (see Section 34.14.4).  | If IA32_VMX_MISC[28]  |
|                  |         | 11:3   | Reserved  |   |
|                  |         | 31:12  | MSEG Base (R/W)   |   |
|                  |         | 63:32  | Reserved  |   |
| 9EH              | 158     | IA32_SMBASE  | Base address of the logical processor's SMRAM image (RO, SMM only).   | If IA32_VMX_MISC[15]  |
| C1H              | 193     | IA32_PMC0<br>(PERFCTR0)                                  | General Performance Counter 0 (R/W)   | If CPUID.0AH: EAX[15:8] > 0   |
| C2H              | 194     | IA32_PMC1<br>(PERFCTR1)                                  | General Performance Counter 1 (R/W)   | If CPUID.0AH: EAX[15:8] > 1   |
| C3H              | 195     | IA32_PMC2  | General Performance Counter 2 (R/W)   | If CPUID.0AH: EAX[15:8] > 2   |
| C4H              | 196     | IA32_PMC3  | General Performance Counter 3 (R/W)   | If CPUID.0AH: EAX[15:8] > 3   |
| C5H              | 197     | IA32_PMC4  | General Performance Counter 4 (R/W)   | If CPUID.0AH: EAX[15:8] > 4   |
| C6H              | 198     | IA32_PMC5  | General Performance Counter 5 (R/W)   | If CPUID.0AH: EAX[15:8] > 5   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                            |
|------------------|---------|--|--|------------------------------------|
| Hex              | Decimal |  |  |                                    |
| C7H              | 199     | IA32_PMC6  | General Performance Counter 6 (R/W)  | If CPUID.0AH: EAX[15:8] > 6        |
| C8H              | 200     | IA32_PMC7  | General Performance Counter 7 (R/W)  | If CPUID.0AH: EAX[15:8] > 7        |
| E7H              | 231     | IA32_MPERF   | TSC Frequency Clock Counter (R/Write to clear)   | If CPUID.06H: ECX[0] = 1           |
|                  |         | 63:0   | CO_MCNT: CO TSC Frequency Clock Count Increments at fixed interval (relative to TSC freq.) when the logical processor is in CO. Cleared upon overflow / wrap-around of IA32_APERF.                 |                                    |
| E8H              | 232     | IA32_APERF   | Actual Performance Clock Counter (R/Write to clear)  | If CPUID.06H: ECX[0] = 1           |
|                  |         | 63:0   | CO_ACNT: CO Actual Frequency Clock Count Accumulates core clock counts at the coordinated clock frequency, when the logical processor is in CO. Cleared upon overflow / wrap-around of IA32_MPERF. |                                    |
| FEH              | 254     | IA32_MTRRCAP<br>(MTRRcap)                                | MTRR Capability (RO)<br>See Section 11.11.2.1,<br>"IA32_MTRR_DEF_TYPE MSR."  | 06_01H                             |
|                  |         | 7:0  | VCNT: The number of variable memory type ranges in the processor.  |                                    |
|                  |         | 8  | Fixed range MTRRs are supported when set.  |                                    |
|                  |         | 9  | Reserved   |                                    |
|                  |         | 10   | WC Supported when set.   |                                    |
|                  |         | 11   | SMRR Supported when set.   |                                    |
|                  |         | 63:12  | Reserved   |                                    |
| 174H             | 372     | IA32_SYSENTER_CS   | SYSENTER_CS_MSR (R/W)  | 06_01H                             |
|                  |         | 15:0   | CS Selector.   |                                    |
|                  |         | 31:16  | Not used.  | Can be read and written.           |
|                  |         | 63:32  | Not used.  | Writes ignored; reads return zero. |
| 175H             | 373     | IA32_SYSENTER_ESP  | SYSENTER_ESP_MSR (R/W)   | 06_01H                             |
| 176H             | 374     | IA32_SYSENTER_EIP  | SYSENTER_EIP_MSR (R/W)   | 06_01H                             |
| 179H             | 377     | IA32_MCG_CAP (MCG_CAP)                                   | Global Machine Check Capability (RO)   | 06_01H                             |
|                  |         | 7:0  | Count: Number of reporting banks.  |                                    |
|                  |         | 8  | MCG_CTL_P: IA32_MCG_CTL is present if this bit is set.   |                                    |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                         |
|------------------|---------|--|---|---------------------------------|
| Hex              | Decimal |  |   |                                 |
|                  |         | 9  | MCG_EXT_P: Extended machine check state registers are present if this bit is set.   |                                 |
|                  |         | 10   | MCP_CMCI_P: Support for corrected MC error event is present.  | 06_01H                          |
|                  |         | 11   | MCG_TES_P: Threshold-based error status register are present if this bit is set.  |                                 |
|                  |         | 15:12  | Reserved  |                                 |
|                  |         | 23:16  | MCG_EXT_CNT: Number of extended machine check state registers present.  |                                 |
|                  |         | 24   | MCG_SER_P: The processor supports software error recovery if this bit is set.   |                                 |
|                  |         | 25   | Reserved  |                                 |
|                  |         | 26   | MCG_ELOG_P: Indicates that the processor allows platform firmware to be invoked when an error is detected so that it may provide additional platform specific information in an ACPI format "Generic Error Data Entry" that augments the data included in machine check bank registers. | 06_3EH                          |
|                  |         | 27   | MCG_LMCE_P: Indicates that the processor supports extended state in IA32_MCG_STATUS and associated MSR necessary to configure Local Machine Check Exception (LMCE).   | 06_3EH                          |
|                  | 63:28   | Reserved   |   |                                 |
| 17AH             | 378     | IA32_MCG_STATUS (MCG_STATUS)                             | Global Machine Check Status (R/W0)  | 06_01H                          |
|                  |         | 0  | RIPV. Restart IP valid.   | 06_01H                          |
|                  |         | 1  | EIPV. Error IP valid.   | 06_01H                          |
|                  |         | 2  | MCIP. Machine check in progress.  | 06_01H                          |
|                  |         | 3  | LMCE_S  | If IA32_MCG_CAP.LMCE_P[2:7] = 1 |
|                  | 63:4    | Reserved   |   |                                 |
| 17BH             | 379     | IA32_MCG_CTL (MCG_CTL)                                   | Global Machine Check Control (R/W)  | If IA32_MCG_CAP.CTL_P[8] = 1    |
| 180H-185H        | 384-389 | Reserved   |   | 06_0EH <sup>1</sup>             |
| 186H             | 390     | IA32_PERFEVTSELO (PERFEVTSELO)                           | Performance Event Select Register 0 (R/W)   | If CPUID.0AH: EAX[15:8] > 0     |
|                  |         | 7:0  | Event Select: Selects a performance event logic unit.   |                                 |



Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                     |
|------------------|---------|--|---|-----------------------------|
| Hex              | Decimal |  |   |                             |
|                  |         | 15:8   | UMask: Qualifies the microarchitectural condition to detect on the selected event logic.  |                             |
|                  |         | 16   | USR: Counts while in privilege level is not ring 0.   |                             |
|                  |         | 17   | OS: Counts while in privilege level is ring 0.  |                             |
|                  |         | 18   | Edge: Enables edge detection if set.  |                             |
|                  |         | 19   | PC: Enables pin control.  |                             |
|                  |         | 20   | INT: Enables interrupt on counter overflow.   |                             |
|                  |         | 21   | AnyThread: When set to 1, it enables counting the associated event conditions occurring across all logical processors sharing a processor core. When set to 0, the counter only increments the associated event conditions occurring in the logical processor which programmed the MSR. |                             |
|                  |         | 22   | EN: Enables the corresponding performance counter to commence counting when this bit is set.  |                             |
|                  |         | 23   | INV: Invert the CMASK.  |                             |
|                  |         | 31:24  | CMASK: When CMASK is not zero, the corresponding performance counter increments each cycle if the event count is greater than or equal to the CMASK.  |                             |
|                  |         | 63:32  | Reserved  |                             |
| 187H             | 391     | IA32_PERFEVTSEL1<br>(PERFEVTSEL1)                        | Performance Event Select Register 1 (R/W)   | If CPUID.0AH: EAX[15:8] > 1 |
| 188H             | 392     | IA32_PERFEVTSEL2   | Performance Event Select Register 2 (R/W)   | If CPUID.0AH: EAX[15:8] > 2 |
| 189H             | 393     | IA32_PERFEVTSEL3   | Performance Event Select Register 3 (R/W)   | If CPUID.0AH: EAX[15:8] > 3 |
| 18AH-197H        | 394-407 | Reserved   |   | 06_0EH <sup>2</sup>         |
| 198H             | 408     | IA32_PERF_STATUS   | Current Performance Status (RO)<br>See Section 14.1.1, "Software Interface For Initiating Performance State Transitions".   | 0F_03H                      |
|                  |         | 15:0   | Current performance State Value.  |                             |
|                  |         | 63:16  | Reserved  |                             |
| 199H             | 409     | IA32_PERF_CTL  | Performance Control MSR (R/W)<br>Software makes a request for a new Performance state (P-State) by writing this MSR. See Section 14.1.1, "Software Interface For Initiating Performance State Transitions".   | 0F_03H                      |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                  |
|------------------|---------|--|--|--------------------------|
| Hex              | Decimal |  |  |                          |
|                  |         | 15:0   | Target performance State Value.  |                          |
|                  |         | 31:16  | Reserved   |                          |
|                  |         | 32   | IDA Engage (R/W)<br>When set to 1: disengages IDA.   | 06_0FH (Mobile only)     |
|                  |         | 63:33  | Reserved   |                          |
| 19AH             | 410     | IA32_CLOCK_MODULATION                                    | Clock Modulation Control (R/W)<br>See Section 14.7.3, "Software Controlled Clock Modulation."  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 0  | Extended On-Demand Clock Modulation Duty Cycle.  | If CPUID.06H:EAX[5] = 1  |
|                  |         | 3:1  | On-Demand Clock Modulation Duty Cycle: Specific encoded values for target duty cycle modulation.   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 4  | On-Demand Clock Modulation Enable: Set 1 to enable modulation.   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 63:5   | Reserved   |                          |
| 19BH             | 411     | IA32_THERM_INTERRUPT                                     | Thermal Interrupt Control (R/W)<br>Enables and disables the generation of an interrupt on temperature transitions detected with the processor's thermal sensors and thermal monitor.<br>See Section 14.7.2, "Thermal Monitor." | If CPUID.01H:EDX[22] = 1 |
|                  |         | 0  | High-Temperature Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 1  | Low-Temperature Interrupt Enable   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 2  | PROCHOT# Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 3  | FORCEPR# Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 4  | Critical Temperature Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 7:5  | Reserved   |                          |
|                  |         | 14:8   | Threshold #1 Value   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 15   | Threshold #1 Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 22:16  | Threshold #2 Value   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 23   | Threshold #2 Interrupt Enable  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 24   | Power Limit Notification Enable  | If CPUID.06H:EAX[4] = 1  |
| 19CH             | 412     | IA32_THERM_STATUS  | Thermal Status Information (RO)<br>Contains status information about the processor's thermal sensor and automatic thermal monitoring facilities.<br>See Section 14.7.2, "Thermal Monitor".                                     | If CPUID.01H:EDX[22] = 1 |
|                  |         | 0  | Thermal Status (RO)  | If CPUID.01H:EDX[22] = 1 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                  |
|------------------|---------|--|--|--------------------------|
| Hex              | Decimal |  |  |                          |
|                  |         | 1  | Thermal Status Log (R/W)   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 2  | PROCHOT # or FORCEPR# event (RO)   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 3  | PROCHOT # or FORCEPR# log (R/WCO)  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 4  | Critical Temperature Status (RO)   | If CPUID.01H:EDX[22] = 1 |
|                  |         | 5  | Critical Temperature Status log (R/WCO)  | If CPUID.01H:EDX[22] = 1 |
|                  |         | 6  | Thermal Threshold #1 Status (RO)   | If CPUID.01H:ECX[8] = 1  |
|                  |         | 7  | Thermal Threshold #1 log (R/WCO)   | If CPUID.01H:ECX[8] = 1  |
|                  |         | 8  | Thermal Threshold #2 Status (RO)   | If CPUID.01H:ECX[8] = 1  |
|                  |         | 9  | Thermal Threshold #2 log (R/WCO)   | If CPUID.01H:ECX[8] = 1  |
|                  |         | 10   | Power Limitation Status (RO)   | If CPUID.06H:EAX[4] = 1  |
|                  |         | 11   | Power Limitation log (R/WCO)   | If CPUID.06H:EAX[4] = 1  |
|                  |         | 12   | Current Limit Status (RO)  | If CPUID.06H:EAX[7] = 1  |
|                  |         | 13   | Current Limit log (R/WCO)  | If CPUID.06H:EAX[7] = 1  |
|                  |         | 14   | Cross Domain Limit Status (RO)   | If CPUID.06H:EAX[7] = 1  |
|                  |         | 15   | Cross Domain Limit log (R/WCO)   | If CPUID.06H:EAX[7] = 1  |
|                  |         | 22:16  | Digital Readout (RO)   | If CPUID.06H:EAX[0] = 1  |
|                  |         | 26:23  | Reserved   |                          |
|                  |         | 30:27  | Resolution in Degrees Celsius (RO)   | If CPUID.06H:EAX[0] = 1  |
|                  |         | 31   | Reading Valid (RO)   | If CPUID.06H:EAX[0] = 1  |
|                  |         | 63:32  | Reserved   |                          |
| 1A0H             | 416     | IA32_MISC_ENABLE   | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.   |                          |
|                  |         | 0  | Fast-Strings Enable<br>When set, the fast-strings feature (for REP MOVS and REP STORS) is enabled (default).<br>When clear, fast-strings are disabled. | OF_OH                    |
|                  |         | 2:1  | Reserved   |                          |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                 |
|------------------|---------|--|---|-------------------------|
| Hex              | Decimal |  |   |                         |
|                  |         | 3  | Automatic Thermal Control Circuit Enable (R/W)<br>1 = Setting this bit enables the thermal control circuit (TCC) portion of the Intel Thermal Monitor feature. This allows the processor to automatically reduce power consumption in response to TCC activation.<br>0 = Disabled.<br>Note: In some products clearing this bit might be ignored in critical thermal conditions, and TM1, TM2 and adaptive thermal throttling will still be activated.<br>The default value of this field varies with product . See respective tables where default value is listed. | 0F_0H                   |
|                  |         | 6:4  | Reserved  |                         |
|                  |         | 7  | Performance Monitoring Available (R)<br>1 = Performance monitoring enabled.<br>0 = Performance monitoring disabled.   | 0F_0H                   |
|                  |         | 10:8   | Reserved  |                         |
|                  |         | 11   | Branch Trace Storage Unavailable (RO)<br>1 = Processor doesn't support branch trace storage (BTS).<br>0 = BTS is supported.   | 0F_0H                   |
|                  |         | 12   | Processor Event Based Sampling (PEBS) Unavailable (RO)<br>1 = PEBS is not supported.<br>0 = PEBS is supported.  | 06_0FH                  |
|                  |         | 15:13  | Reserved  |                         |
|                  |         | 16   | Enhanced Intel SpeedStep Technology Enable (R/W)<br>0= Enhanced Intel SpeedStep Technology disabled.<br>1 = Enhanced Intel SpeedStep Technology enabled.  | If CPUID.01H: ECX[7] =1 |
|                  |         | 17   | Reserved  |                         |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                  |
|------------------|---------|--|---|--------------------------|
| Hex              | Decimal |  |   |                          |
|                  |         | 18   | <p>ENABLE MONITOR FSM (R/W)</p> <p>When this bit is set to 0, the MONITOR feature flag is not set (CPUID.01H:ECX[bit 3] = 0). This indicates that MONITOR/MWAIT are not supported.</p> <p>Software attempts to execute MONITOR/MWAIT will cause #UD when this bit is 0.</p> <p>When this bit is set to 1 (default), MONITOR/MWAIT are supported (CPUID.01H:ECX[bit 3] = 1).</p> <p>If the SSE3 feature flag ECX[0] is not set (CPUID.01H:ECX[bit 0] = 0), the OS must not attempt to alter this bit. BIOS must leave it in the default state. Writing this bit when the SSE3 feature flag is set to 0 may generate a #GP exception.</p>   | 0F_03H                   |
|                  |         | 21:19  | Reserved  |                          |
|                  |         | 22   | <p>Limit CPUID Maxval (R/W)</p> <p>When this bit is set to 1, CPUID.00H returns a maximum value in EAX[7:0] of 2.</p> <p>BIOS should contain a setup question that allows users to specify when the installed OS does not support CPUID functions greater than 2.</p> <p>Before setting this bit, BIOS must execute the CPUID.0H and examine the maximum value returned in EAX[7:0]. If the maximum value is greater than 2, this bit is supported.</p> <p>Otherwise, this bit is not supported. Setting this bit when the maximum value is not greater than 2 may generate a #GP exception.</p> <p>Setting this bit may cause unexpected behavior in software that depends on the availability of CPUID leaves greater than 2.</p> | 0F_03H                   |
|                  |         | 23   | <p>xTPR Message Disable (R/W)</p> <p>When set to 1, xTPR messages are disabled. xTPR messages are optional messages that allow the processor to inform the chipset of its priority.</p>   | If CPUID.01H:ECX[14] = 1 |
|                  |         | 33:24  | Reserved  |                          |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |                          | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                        |
|------------------|--------------------------|--|--|--------------------------------|
| Hex              | Decimal                  |  |  |                                |
|                  |                          | 34   | <p>XD Bit Disable (R/W)</p> <p>When set to 1, the Execute Disable Bit feature (XD Bit) is disabled and the XD Bit extended feature flag will be clear (CPUID.80000001H: EDX[20]=0).</p> <p>When set to a 0 (default), the Execute Disable Bit feature (if available) allows the OS to enable PAE paging and take advantage of data only pages.</p> <p>BIOS must not alter the contents of this bit location, if XD bit is not supported. Writing this bit to 1 when the XD Bit extended feature flag is set to 0 may generate a #GP exception.</p> | If CPUID.80000001H:EDX[20] = 1 |
|                  |                          | 63:35  | Reserved   |                                |
| 1B0H             | 432                      | IA32_ENERGY_PERF_BIAS                                    | Performance Energy Bias Hint (R/W)   | If CPUID.6H:ECX[3] = 1         |
|                  |                          | 3:0  | <p>Power Policy Preference:</p> <p>0 indicates preference to highest performance.</p> <p>15 indicates preference to maximize energy saving.</p>  |                                |
|                  |                          | 63:4   | Reserved   |                                |
| 1B1H             | 433                      | IA32_PACKAGE_THERM_STATUS                                | <p>Package Thermal Status Information (RO)</p> <p>Contains status information about the package's thermal sensor.</p> <p>See Section 14.8, "Package Level Thermal Management."</p>   | If CPUID.06H: EAX[6] = 1       |
|                  |                          | 0  | Pkg Thermal Status (RO)  |                                |
|                  |                          | 1  | Pkg Thermal Status Log (R/W)   |                                |
|                  |                          | 2  | Pkg PROCHOT # event (RO)   |                                |
|                  |                          | 3  | Pkg PROCHOT # log (R/WCO)  |                                |
|                  |                          | 4  | Pkg Critical Temperature Status (RO)   |                                |
|                  |                          | 5  | Pkg Critical Temperature Status Log (R/WCO)  |                                |
|                  |                          | 6  | Pkg Thermal Threshold #1 Status (RO)   |                                |
|                  |                          | 7  | Pkg Thermal Threshold #1 log (R/WCO)   |                                |
|                  |                          | 8  | Pkg Thermal Threshold #2 Status (RO)   |                                |
|                  |                          | 9  | Pkg Thermal Threshold #1 log (R/WCO)   |                                |
|                  |                          | 10   | Pkg Power Limitation Status (RO)   |                                |
|                  |                          | 11   | Pkg Power Limitation log (R/WCO)   |                                |
|                  |                          | 15:12  | Reserved   |                                |
| 22:16            | Pkg Digital Readout (RO) |  |  |                                |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |  | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|--|--|--|---|
| Hex              | Decimal  |  |  |   |
|                  |  | 63:23  | Reserved   |   |
| 1B2H             | 434  | IA32_PACKAGE_THERM_INTERRUPT                             | Pkg Thermal Interrupt Control (R/W)<br>Enables and disables the generation of an interrupt on temperature transitions detected with the package's thermal sensor.<br>See Section 14.8, "Package Level Thermal Management." | If CPUID.06H: EAX[6] = 1                        |
|                  |  | 0  | Pkg High-Temperature Interrupt Enable  |   |
|                  |  | 1  | Pkg Low-Temperature Interrupt Enable   |   |
|                  |  | 2  | Pkg PROCHOT# Interrupt Enable  |   |
|                  |  | 3  | Reserved   |   |
|                  |  | 4  | Pkg Overheat Interrupt Enable  |   |
|                  |  | 7:5  | Reserved   |   |
|                  |  | 14:8   | Pkg Threshold #1 Value   |   |
|                  |  | 15   | Pkg Threshold #1 Interrupt Enable  |   |
|                  |  | 22:16  | Pkg Threshold #2 Value   |   |
|                  |  | 23   | Pkg Threshold #2 Interrupt Enable  |   |
|                  |  | 24   | Pkg Power Limit Notification Enable  |   |
|                  |  | 63:25  | Reserved   |   |
|                  |  | 1D9H   | 473  | IA32_DEBUGCTL<br>(MSR_DEBUGCTLA, MSR_DEBUGCTLB) |
| 0                | LBR: Setting this bit to 1 enables the processor to record a running trace of the most recent branches taken by the processor in the LBR stack.                          |  |  | 06_01H  |
| 1                | BTF: Setting this bit to 1 enables the processor to treat EFLAGS.TF as single-step on branches instead of single-step on instructions.                                   |  |  | 06_01H  |
| 5:2              | Reserved   |  |  |   |
| 6                | TR: Setting this bit to 1 enables branch trace messages to be sent.  |  |  | 06_0EH  |
| 7                | BTS: Setting this bit enables branch trace messages (BTMs) to be logged in a BTS buffer.   |  |  | 06_0EH  |
| 8                | BTINT: When clear, BTMs are logged in a BTS buffer in circular fashion. When this bit is set, an interrupt is generated by the BTS facility when the BTS buffer is full. |  |  | 06_0EH  |
| 9                | 1: BTS_OFF_OS: When set, BTS or BTM is skipped if CPL = 0.   |  |  | 06_0FH  |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment  |
|------------------|---------|--|--|--|
| Hex              | Decimal |  |  |  |
|                  |         | 10   | BTS_OFF_USR: When set, BTS or BTM is skipped if CPL > 0.   | 06_0FH   |
|                  |         | 11   | FREEZE_LBRS_ON_PMI: When set, the LBR stack is frozen on a PMI request.  | If CPUID.01H: ECX[15] = 1 && CPUID.0AH: EAX[7:0] > 1 |
|                  |         | 12   | FREEZE_PERFMON_ON_PMI: When set, each ENABLE bit of the global counter control MSR are frozen (address 38FH) on a PMI request. | If CPUID.01H: ECX[15] = 1 && CPUID.0AH: EAX[7:0] > 1 |
|                  |         | 13   | ENABLE_UNCORE_PMI: When set, enables the logical processor to receive and generate PMI on behalf of the uncore.                | 06_1AH   |
|                  |         | 14   | FREEZE_WHILE_SMM: When set, freezes perfmon and trace messages while in SMM.   | If IA32_PERF_CAPABILITIES[12] = 1                    |
|                  |         | 15   | RTM_DEBUG: When set, enables DR7 debug bit on XBEGIN.  | If (CPUID.(EAX=07H, ECX=0):EBX[11] = 1)              |
|                  |         | 63:16  | Reserved   |  |
| 1F2H             | 498     | IA32_SMRR_PHYSBASE                                       | SMRR Base Address (Writeable only in SMM)<br>Base address of SMM memory range.   | If IA32_MTRRCAP.SMRR[11] = 1                         |
|                  |         | 7:0  | Type. Specifies memory type of the range.  |  |
|                  |         | 11:8   | Reserved   |  |
|                  |         | 31:12  | PhysBase<br>SMRR physical Base Address.  |  |
|                  |         | 63:32  | Reserved   |  |
| 1F3H             | 499     | IA32_SMRR_PHYSMASK                                       | SMRR Range Mask (Writeable only in SMM)<br>Range Mask of SMM memory range.   | If IA32_MTRRCAP[SMRR] = 1                            |
|                  |         | 10:0   | Reserved   |  |
|                  |         | 11   | Valid<br>Enable range mask.  |  |
|                  |         | 31:12  | PhysMask<br>SMRR address range mask.   |  |
|                  |         | 63:32  | Reserved   |  |
| 1F8H             | 504     | IA32_PLATFORM_DCA_CAP                                    | DCA Capability (R)   | If CPUID.01H: ECX[18] = 1                            |
| 1F9H             | 505     | IA32_CPU_DCA_CAP   | If set, CPU supports Prefetch-Hint type.   | If CPUID.01H: ECX[18] = 1                            |
| 1FAH             | 506     | IA32_DCA_0_CAP   | DCA type 0 Status and Control register.  | If CPUID.01H: ECX[18] = 1                            |
|                  |         | 0  | DCA_ACTIVE: Set by HW when DCA is fuse-enabled and no defeatures are set.  |  |
|                  |         | 2:1  | TRANSACTION  |  |
|                  |         | 6:3  | DCA_TYPE   |  |



Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                           |
|------------------|---------|--|--|-----------------------------------|
| Hex              | Decimal |  |  |                                   |
|                  |         | 10:7   | DCA_QUEUE_SIZE   |                                   |
|                  |         | 12:11  | Reserved   |                                   |
|                  |         | 16:13  | DCA_DELAY: Writes will update the register but have no HW side-effect. |                                   |
|                  |         | 23:17  | Reserved   |                                   |
|                  |         | 24   | SW_BLOCK: SW can request DCA block by setting this bit.                |                                   |
|                  |         | 25   | Reserved   |                                   |
|                  |         | 26   | HW_BLOCK: Set when DCA is blocked by HW (e.g. CRO.CD = 1).             |                                   |
|                  |         | 31:27  | Reserved   |                                   |
| 200H             | 512     | IA32_MTRR_PHYSBASE0<br>(MTRRphysBase0)                   | See Section 11.11.2.3, "Variable Range MTRRs."                         | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 201H             | 513     | IA32_MTRR_PHYSMASK0                                      | MTRRphysMask0  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 202H             | 514     | IA32_MTRR_PHYSBASE1                                      | MTRRphysBase1  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 203H             | 515     | IA32_MTRR_PHYSMASK1                                      | MTRRphysMask1  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 204H             | 516     | IA32_MTRR_PHYSBASE2                                      | MTRRphysBase2  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 205H             | 517     | IA32_MTRR_PHYSMASK2                                      | MTRRphysMask2  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 206H             | 518     | IA32_MTRR_PHYSBASE3                                      | MTRRphysBase3  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 207H             | 519     | IA32_MTRR_PHYSMASK3                                      | MTRRphysMask3  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 208H             | 520     | IA32_MTRR_PHYSBASE4                                      | MTRRphysBase4  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 209H             | 521     | IA32_MTRR_PHYSMASK4                                      | MTRRphysMask4  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20AH             | 522     | IA32_MTRR_PHYSBASE5                                      | MTRRphysBase5  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20BH             | 523     | IA32_MTRR_PHYSMASK5                                      | MTRRphysMask5  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20CH             | 524     | IA32_MTRR_PHYSBASE6                                      | MTRRphysBase6  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20DH             | 525     | IA32_MTRR_PHYSMASK6                                      | MTRRphysMask6  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20EH             | 526     | IA32_MTRR_PHYSBASE7                                      | MTRRphysBase7  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 20FH             | 527     | IA32_MTRR_PHYSMASK7                                      | MTRRphysMask7  | If CPUID.01H:<br>EDX.MTRR[12] = 1 |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                         | Comment                           |
|------------------|---------|--|---|-----------------------------------|
| Hex              | Decimal |  |   |                                   |
| 210H             | 528     | IA32_MTRR_PHYSBASE8                                      | MTRRphysBase8                               | if IA32_MTRRCAP[7:0] > 8          |
| 211H             | 529     | IA32_MTRR_PHYSMASK8                                      | MTRRphysMask8                               | if IA32_MTRRCAP[7:0] > 8          |
| 212H             | 530     | IA32_MTRR_PHYSBASE9                                      | MTRRphysBase9                               | if IA32_MTRRCAP[7:0] > 9          |
| 213H             | 531     | IA32_MTRR_PHYSMASK9                                      | MTRRphysMask9                               | if IA32_MTRRCAP[7:0] > 9          |
| 250H             | 592     | IA32_MTRR_FIX64K_00000                                   | MTRRfix64K_00000                            | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 258H             | 600     | IA32_MTRR_FIX16K_80000                                   | MTRRfix16K_80000                            | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 259H             | 601     | IA32_MTRR_FIX16K_A0000                                   | MTRRfix16K_A0000                            | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 268H             | 616     | IA32_MTRR_FIX4K_C0000<br>(MTRRfix4K_C0000 )              | See Section 11.11.2.2, "Fixed Range MTRRs." | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 269H             | 617     | IA32_MTRR_FIX4K_C8000                                    | MTRRfix4K_C8000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26AH             | 618     | IA32_MTRR_FIX4K_D0000                                    | MTRRfix4K_D0000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26BH             | 619     | IA32_MTRR_FIX4K_D8000                                    | MTRRfix4K_D8000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26CH             | 620     | IA32_MTRR_FIX4K_E0000                                    | MTRRfix4K_E0000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26DH             | 621     | IA32_MTRR_FIX4K_E8000                                    | MTRRfix4K_E8000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26EH             | 622     | IA32_MTRR_FIX4K_F0000                                    | MTRRfix4K_F0000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 26FH             | 623     | IA32_MTRR_FIX4K_F8000                                    | MTRRfix4K_F8000                             | If CPUID.01H:<br>EDX.MTRR[12] = 1 |
| 277H             | 631     | IA32_PAT   | IA32_PAT (R/W)                              | If CPUID.01H:<br>EDX.MTRR[16] = 1 |
|                  |         | 2:0  | PA0   |                                   |
|                  |         | 7:3  | Reserved                                    |                                   |
|                  |         | 10:8   | PA1   |                                   |
|                  |         | 15:11  | Reserved                                    |                                   |
|                  |         | 18:16  | PA2   |                                   |
|                  |         | 23:19  | Reserved                                    |                                   |
|                  |         | 26:24  | PA3   |                                   |
|                  |         | 31:27  | Reserved                                    |                                   |
|                  |         | 34:32  | PA4   |                                   |
|                  |         | 39:35  | Reserved                                    |                                   |
|                  |         | 42:40  | PA5   |                                   |
|                  |         | 47:43  | Reserved                                    |                                   |
| 50:48            | PA6     |  |   |                                   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
|                  |         | 55:51  | Reserved   |   |
|                  |         | 58:56  | PA7  |   |
|                  |         | 63:59  | Reserved   |   |
| 280H             | 640     | IA32_MCO_CTL2  | MSR to enable/disable CMCI capability for bank 0. (R/W)<br>See Section 15.3.2.5, "IA32_MCi_CTL2 MSRs". | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 0  |
|                  |         | 14:0   | Corrected error count threshold.   |   |
|                  |         | 29:15  | Reserved   |   |
|                  |         | 30   | CMCI_EN  |   |
|                  |         | 63:31  | Reserved   |   |
| 281H             | 641     | IA32_MC1_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 1  |
| 282H             | 642     | IA32_MC2_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 2  |
| 283H             | 643     | IA32_MC3_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 3  |
| 284H             | 644     | IA32_MC4_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 4  |
| 285H             | 645     | IA32_MC5_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 5  |
| 286H             | 646     | IA32_MC6_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 6  |
| 287H             | 647     | IA32_MC7_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 7  |
| 288H             | 648     | IA32_MC8_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 8  |
| 289H             | 649     | IA32_MC9_CTL2  | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 9  |
| 28AH             | 650     | IA32_MC10_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 10 |
| 28BH             | 651     | IA32_MC11_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 11 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                 | Comment  |
|------------------|---------|--|-------------------------------------|--|
| Hex              | Decimal |  |                                     |  |
| 28CH             | 652     | IA32_MC12_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 12 |
| 28DH             | 653     | IA32_MC13_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 13 |
| 28EH             | 654     | IA32_MC14_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 14 |
| 28FH             | 655     | IA32_MC15_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 15 |
| 290H             | 656     | IA32_MC16_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 16 |
| 291H             | 657     | IA32_MC17_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 17 |
| 292H             | 658     | IA32_MC18_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 18 |
| 293H             | 659     | IA32_MC19_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 19 |
| 294H             | 660     | IA32_MC20_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 20 |
| 295H             | 661     | IA32_MC21_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 21 |
| 296H             | 662     | IA32_MC22_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 22 |
| 297H             | 663     | IA32_MC23_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 23 |
| 298H             | 664     | IA32_MC24_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 24 |
| 299H             | 665     | IA32_MC25_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 25 |
| 29AH             | 666     | IA32_MC26_CTL2   | (R/W) Same fields as IA32_MCO_CTL2. | If IA32_MCG_CAP[10] = 1<br>&& IA32_MCG_CAP[7:0] > 26 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
| 29BH             | 667     | IA32_MC27_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 27 |
| 29CH             | 668     | IA32_MC28_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 28 |
| 29DH             | 669     | IA32_MC29_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 29 |
| 29EH             | 670     | IA32_MC30_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 30 |
| 29FH             | 671     | IA32_MC31_CTL2   | (R/W) Same fields as IA32_MCO_CTL2.  | If IA32_MCG_CAP[10] = 1 && IA32_MCG_CAP[7:0] > 31 |
| 2FFH             | 767     | IA32_MTRR_DEF_TYPE                                       | MTRRdefType (R/W)  | If CPUID.01H: EDX.MTRR[12] = 1                    |
|                  |         | 2:0  | Default Memory Type  |   |
|                  |         | 9:3  | Reserved   |   |
|                  |         | 10   | Fixed Range MTRR Enable  |   |
|                  |         | 11   | MTRR Enable  |   |
|                  |         | 63:12  | Reserved   |   |
| 309H             | 777     | IA32_FIXED_CTR0<br>(MSR_PERF_FIXED_CTR0)                 | Fixed-Function Performance Counter 0<br>(R/W): Counts Instr_Retired.Any.             | If CPUID.0AH: EDX[4:0] > 0                        |
| 30AH             | 778     | IA32_FIXED_CTR1<br>(MSR_PERF_FIXED_CTR1)                 | Fixed-Function Performance Counter 1<br>(R/W): Counts CPU_CLK_Unhalted.Core.         | If CPUID.0AH: EDX[4:0] > 1                        |
| 30BH             | 779     | IA32_FIXED_CTR2<br>(MSR_PERF_FIXED_CTR2)                 | Fixed-Function Performance Counter 2<br>(R/W): Counts CPU_CLK_Unhalted.Ref.          | If CPUID.0AH: EDX[4:0] > 2                        |
| 345H             | 837     | IA32_PERF_CAPABILITIES                                   | Read Only MSR that enumerates the existence of performance monitoring features. (RO) | If CPUID.01H: ECX[15] = 1                         |
|                  |         | 5:0  | LBR format   |   |
|                  |         | 6  | PEBS Trap  |   |
|                  |         | 7  | PEBSSaveArchRegs   |   |
|                  |         | 11:8   | PEBS Record Format   |   |
|                  |         | 12   | 1: Freeze while SMM is supported.  |   |
|                  |         | 13   | 1: Full width of counter writable via IA32_A_PMCx.                                   |   |
|                  |         | 63:14  | Reserved   |   |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                    |
|------------------|---------|--|---|----------------------------|
| Hex              | Decimal |  |   |                            |
| 38DH             | 909     | IA32_FIXED_CTR_CTRL                                      | Fixed-Function Performance Counter Control (R/W)<br><br>Counter increments while the results of ANDing respective enable bit in IA32_PERF_GLOBAL_CTRL with the corresponding OS or USR bits in this MSR is true.  | If CPUID.0AH: EAX[7:0] > 1 |
|                  |         | 0  | EN0_OS: Enable Fixed Counter 0 to count while CPL = 0.  |                            |
|                  |         | 1  | EN0_Usr: Enable Fixed Counter 0 to count while CPL > 0.   |                            |
|                  |         | 2  | AnyThread: When set to 1, it enables counting the associated event conditions occurring across all logical processors sharing a processor core. When set to 0, the counter only increments the associated event conditions occurring in the logical processor which programmed the MSR. | If CPUID.0AH: EAX[7:0] > 2 |
|                  |         | 3  | EN0_PMI: Enable PMI when fixed counter 0 overflows.   |                            |
|                  |         | 4  | EN1_OS: Enable Fixed Counter 1 to count while CPL = 0.  |                            |
|                  |         | 5  | EN1_Usr: Enable Fixed Counter 1 to count while CPL > 0.   |                            |
|                  |         | 6  | AnyThread: When set to 1, it enables counting the associated event conditions occurring across all logical processors sharing a processor core. When set to 0, the counter only increments the associated event conditions occurring in the logical processor which programmed the MSR. | If CPUID.0AH: EAX[7:0] > 2 |
|                  |         | 7  | EN1_PMI: Enable PMI when fixed counter 1 overflows.   |                            |
|                  |         | 8  | EN2_OS: Enable Fixed Counter 2 to count while CPL = 0.  |                            |
|                  |         | 9  | EN2_Usr: Enable Fixed Counter 2 to count while CPL > 0.   |                            |
|                  |         | 10   | AnyThread: When set to 1, it enables counting the associated event conditions occurring across all logical processors sharing a processor core. When set to 0, the counter only increments the associated event conditions occurring in the logical processor which programmed the MSR. | If CPUID.0AH: EAX[7:0] > 2 |
|                  |         | 11   | EN2_PMI: Enable PMI when fixed counter 2 overflows.   |                            |
|                  | 63:12   | Reserved   |   |                            |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |  | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|--|--|--|---|
| Hex              | Decimal  |  |  |   |
| 38EH             | 910  | IA32_PERF_GLOBAL_STATUS                                  | Global Performance Counter Status (RO)   | If CPUID.0AH: EAX[7:0] > 0  |
|                  |  | 0  | Ovf_PMC0: Overflow status of IA32_PMC0.  | If CPUID.0AH: EAX[15:8] > 0                                       |
|                  |  | 1  | Ovf_PMC1: Overflow status of IA32_PMC1.  | If CPUID.0AH: EAX[15:8] > 1                                       |
|                  |  | 2  | Ovf_PMC2: Overflow status of IA32_PMC2.  | If CPUID.0AH: EAX[15:8] > 2                                       |
|                  |  | 3  | Ovf_PMC3: Overflow status of IA32_PMC3.  | If CPUID.0AH: EAX[15:8] > 3                                       |
|                  |  | 31:4   | Reserved   |   |
|                  |  | 32   | Ovf_FixedCtr0: Overflow status of IA32_FIXED_CTR0.   | If CPUID.0AH: EAX[7:0] > 1  |
|                  |  | 33   | Ovf_FixedCtr1: Overflow status of IA32_FIXED_CTR1.   | If CPUID.0AH: EAX[7:0] > 1  |
|                  |  | 34   | Ovf_FixedCtr2: Overflow status of IA32_FIXED_CTR2.   | If CPUID.0AH: EAX[7:0] > 1  |
|                  |  | 54:35  | Reserved   |   |
|                  |  | 55   | Trace_ToPA_PMI: A PMI occurred due to a ToPA entry memory buffer that was completely filled.   | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1) && IA32_RTIT_CTL.ToPA = 1 |
|                  |  | 57:56  | Reserved   |   |
|                  |  | 58   | LBR_Frz. LBRs are frozen due to: <ul style="list-style-type: none"> <li>▪ IA32_DEBUGCTL.FREEZE_LBR_ON_PMI=1.</li> <li>▪ The LBR stack overflowed.</li> </ul>   | If CPUID.0AH: EAX[7:0] > 3  |
|                  |  | 59   | CTR_Frz. Performance counters in the core PMU are frozen due to: <ul style="list-style-type: none"> <li>▪ IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI=1.</li> <li>▪ One or more core PMU counters overflowed.</li> </ul> | If CPUID.0AH: EAX[7:0] > 3  |
|                  |  | 60   | ASCI: Data in the performance counters in the core PMU may include contributions from the direct or indirect operation Intel SGX to protect an enclave.  | If CPUID.(EAX=07H, ECX=0):EBX[2] = 1                              |
|                  |  | 61   | Ovf_Uncore: Uncore counter overflow status.  | If CPUID.0AH: EAX[7:0] > 2  |
| 62               | OvfBuf: DS SAVE area Buffer overflow status.         | If CPUID.0AH: EAX[7:0] > 0                               |  |   |
| 63               | CondChgd: Status bits of this register have changed. | If CPUID.0AH: EAX[7:0] > 0                               |  |   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |                              | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|------------------------------|--|--|---|
| Hex              | Decimal                      |  |  |   |
| 38FH             | 911                          | IA32_PERF_GLOBAL_CTRL                                    | Global Performance Counter Control (R/W)<br>Counter increments while the result of ANDing the respective enable bit in this MSR with the corresponding OS or USR bits in the general-purpose or fixed counter control MSR is true. | If CPUID.0AH: EAX[7:0] > 0  |
|                  |                              | 0  | EN_PMC0  | If CPUID.0AH: EAX[15:8] > 0                                       |
|                  |                              | 1  | EN_PMC1  | If CPUID.0AH: EAX[15:8] > 1                                       |
|                  |                              | 2  | EN_PMC2  | If CPUID.0AH: EAX[15:8] > 2                                       |
|                  |                              | n  | EN_PMCn  | If CPUID.0AH: EAX[15:8] > n                                       |
|                  |                              | 31:n+1   | Reserved   |   |
|                  |                              | 32   | EN_FIXED_CTR0  | If CPUID.0AH: EDX[4:0] > 0  |
|                  |                              | 33   | EN_FIXED_CTR1  | If CPUID.0AH: EDX[4:0] > 1  |
|                  |                              | 34   | EN_FIXED_CTR2  | If CPUID.0AH: EDX[4:0] > 2  |
|                  |                              | 63:35  | Reserved   |   |
| 390H             | 912                          | IA32_PERF_GLOBAL_OVF_CTRL                                | Global Performance Counter Overflow Control (R/W)  | If CPUID.0AH: EAX[7:0] > 0 && CPUID.0AH: EAX[7:0] <= 3            |
|                  |                              | 0  | Set 1 to Clear Ovf_PMC0 bit.   | If CPUID.0AH: EAX[15:8] > 0                                       |
|                  |                              | 1  | Set 1 to Clear Ovf_PMC1 bit.   | If CPUID.0AH: EAX[15:8] > 1                                       |
|                  |                              | 2  | Set 1 to Clear Ovf_PMC2 bit.   | If CPUID.0AH: EAX[15:8] > 2                                       |
|                  |                              | n  | Set 1 to Clear Ovf_PMCn bit.   | If CPUID.0AH: EAX[15:8] > n                                       |
|                  |                              | 31:n   | Reserved   |   |
|                  |                              | 32   | Set 1 to Clear Ovf_FIXED_CTR0 bit.   | If CPUID.0AH: EDX[4:0] > 0  |
|                  |                              | 33   | Set 1 to Clear Ovf_FIXED_CTR1 bit.   | If CPUID.0AH: EDX[4:0] > 1  |
|                  |                              | 34   | Set 1 to Clear Ovf_FIXED_CTR2 bit.   | If CPUID.0AH: EDX[4:0] > 2  |
|                  |                              | 54:35  | Reserved   |   |
|                  |                              | 55   | Set 1 to Clear Trace_ToPA_PMI bit.   | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1) && IA32_RTIT_CTL.ToPA = 1 |
|                  |                              | 60:56  | Reserved   |   |
|                  |                              | 61   | Set 1 to Clear Ovf_Uncore bit.   | 06_2EH  |
|                  |                              | 62   | Set 1 to Clear OvfBuf bit.   | If CPUID.0AH: EAX[7:0] > 0  |
| 63               | Set 1 to clear CondChgd bit. | If CPUID.0AH: EAX[7:0] > 0                               |  |   |



Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |                              | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                                     | Comment  |
|------------------|------------------------------|--|---|--|
| Hex              | Decimal                      |  |   |  |
| 390H             | 912                          | IA32_PERF_GLOBAL_STATUS_RESET                            | Global Performance Counter Overflow Reset Control (R/W) | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 0  | Set 1 to Clear Ovf_PMC0 bit.                            | If CPUID.OAH: EAX[15:8] > 0  |
|                  |                              | 1  | Set 1 to Clear Ovf_PMC1 bit.                            | If CPUID.OAH: EAX[15:8] > 1  |
|                  |                              | 2  | Set 1 to Clear Ovf_PMC2 bit.                            | If CPUID.OAH: EAX[15:8] > 2  |
|                  |                              | n  | Set 1 to Clear Ovf_PMCn bit.                            | If CPUID.OAH: EAX[15:8] > n  |
|                  |                              | 31:n   | Reserved  |  |
|                  |                              | 32   | Set 1 to Clear Ovf_FIXED_CTR0 bit.                      | If CPUID.OAH: EDX[4:0] > 0   |
|                  |                              | 33   | Set 1 to Clear Ovf_FIXED_CTR1 bit.                      | If CPUID.OAH: EDX[4:0] > 1   |
|                  |                              | 34   | Set 1 to Clear Ovf_FIXED_CTR2 bit.                      | If CPUID.OAH: EDX[4:0] > 2   |
|                  |                              | 54:35  | Reserved  |  |
|                  |                              | 55   | Set 1 to Clear Trace_ToPA_PMI bit.                      | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1) && IA32_RTIT_CTL.ToPA[8] = 1 |
|                  |                              | 57:56  | Reserved  |  |
|                  |                              | 58   | Set 1 to Clear LBR_Frz bit.                             | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 59   | Set 1 to Clear CTR_Frz bit.                             | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 58   | Set 1 to Clear ASCII bit.                               | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 61   | Set 1 to Clear Ovf_Uncore bit.                          | O6_2EH   |
|                  |                              | 62   | Set 1 to Clear OvfBuf bit.                              | If CPUID.OAH: EAX[7:0] > 0   |
| 63               | Set 1 to clear CondChgd bit. | If CPUID.OAH: EAX[7:0] > 0                               |   |  |
| 391H             | 913                          | IA32_PERF_GLOBAL_STATUS_SET                              | Global Performance Counter Overflow Set Control (R/W)   | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 0  | Set 1 to cause Ovf_PMC0 = 1.                            | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 1  | Set 1 to cause Ovf_PMC1 = 1.                            | If CPUID.OAH: EAX[15:8] > 1  |
|                  |                              | 2  | Set 1 to cause Ovf_PMC2 = 1.                            | If CPUID.OAH: EAX[15:8] > 2  |
|                  |                              | n  | Set 1 to cause Ovf_PMCn = 1.                            | If CPUID.OAH: EAX[15:8] > n  |
|                  |                              | 31:n   | Reserved  |  |
|                  |                              | 32   | Set 1 to cause Ovf_FIXED_CTR0 = 1.                      | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 33   | Set 1 to cause Ovf_FIXED_CTR1 = 1.                      | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 34   | Set 1 to cause Ovf_FIXED_CTR2 = 1.                      | If CPUID.OAH: EAX[7:0] > 3   |
|                  |                              | 54:35  | Reserved  |  |
|                  |                              | 55   | Set 1 to cause Trace_ToPA_PMI = 1.                      | If CPUID.OAH: EAX[7:0] > 3   |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                                   | Comment                     |
|------------------|---------|--|---|-----------------------------|
| Hex              | Decimal |  |   |                             |
|                  |         | 57:56  | Reserved  |                             |
|                  |         | 58   | Set 1 to cause LBR_Frz = 1.                           | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 59   | Set 1 to cause CTR_Frz = 1.                           | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 58   | Set 1 to cause ASCII = 1.                             | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 61   | Set 1 to cause Ovf_Uncore = 1.                        | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 62   | Set 1 to cause OvfBuf = 1.                            | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 63   | Reserved  |                             |
| 392H             | 914     | IA32_PERF_GLOBAL_INUSE                                   | Indicator that core perfmon interface is in use. (RO) | If CPUID.0AH: EAX[7:0] > 3  |
|                  |         | 0  | IA32_PERFEVTSEL0 in use.                              |                             |
|                  |         | 1  | IA32_PERFEVTSEL1 in use.                              | If CPUID.0AH: EAX[15:8] > 1 |
|                  |         | 2  | IA32_PERFEVTSEL2 in use.                              | If CPUID.0AH: EAX[15:8] > 2 |
|                  |         | n  | IA32_PERFEVTSELn in use.                              | If CPUID.0AH: EAX[15:8] > n |
|                  |         | 31:n+1   | Reserved  |                             |
|                  |         | 32   | IA32_FIXED_CTR0 in use.                               |                             |
|                  |         | 33   | IA32_FIXED_CTR1 in use.                               |                             |
|                  |         | 34   | IA32_FIXED_CTR2 in use.                               |                             |
|                  |         | 62:35  | Reserved or model specific.                           |                             |
|                  |         | 63   | PMI in use.   |                             |
| 3F1H             | 1009    | IA32_PEBS_ENABLE   | PEBS Control (R/W)                                    |                             |
|                  |         | 0  | Enable PEBS on IA32_PMC0.                             | 06_OFH                      |
|                  |         | 3:1  | Reserved or model specific.                           |                             |
|                  |         | 31:4   | Reserved  |                             |
|                  |         | 35:32  | Reserved or model specific.                           |                             |
|                  |         | 63:36  | Reserved  |                             |
| 400H             | 1024    | IA32_MCO_CTL   | MCO_CTL   | If IA32_MCG_CAP.CNT > 0     |
| 401H             | 1025    | IA32_MCO_STATUS  | MCO_STATUS  | If IA32_MCG_CAP.CNT > 0     |
| 402H             | 1026    | IA32_MCO_ADDR <sup>1</sup>                               | MCO_ADDR  | If IA32_MCG_CAP.CNT > 0     |
| 403H             | 1027    | IA32_MCO_MISC  | MCO_MISC  | If IA32_MCG_CAP.CNT > 0     |
| 404H             | 1028    | IA32_MC1_CTL   | MC1_CTL   | If IA32_MCG_CAP.CNT > 1     |
| 405H             | 1029    | IA32_MC1_STATUS  | MC1_STATUS  | If IA32_MCG_CAP.CNT > 1     |
| 406H             | 1030    | IA32_MC1_ADDR <sup>2</sup>                               | MC1_ADDR  | If IA32_MCG_CAP.CNT > 1     |
| 407H             | 1031    | IA32_MC1_MISC  | MC1_MISC  | If IA32_MCG_CAP.CNT > 1     |
| 408H             | 1032    | IA32_MC2_CTL   | MC2_CTL   | If IA32_MCG_CAP.CNT > 2     |
| 409H             | 1033    | IA32_MC2_STATUS  | MC2_STATUS  | If IA32_MCG_CAP.CNT > 2     |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description | Comment                 |
|------------------|---------|--|---------------------|-------------------------|
| Hex              | Decimal |  |                     |                         |
| 40AH             | 1034    | IA32_MC2_ADDR <sup>1</sup>                               | MC2_ADDR            | If IA32_MCG_CAP.CNT >2  |
| 40BH             | 1035    | IA32_MC2_MISC  | MC2_MISC            | If IA32_MCG_CAP.CNT >2  |
| 40CH             | 1036    | IA32_MC3_CTL   | MC3_CTL             | If IA32_MCG_CAP.CNT >3  |
| 40DH             | 1037    | IA32_MC3_STATUS  | MC3_STATUS          | If IA32_MCG_CAP.CNT >3  |
| 40EH             | 1038    | IA32_MC3_ADDR <sup>1</sup>                               | MC3_ADDR            | If IA32_MCG_CAP.CNT >3  |
| 40FH             | 1039    | IA32_MC3_MISC  | MC3_MISC            | If IA32_MCG_CAP.CNT >3  |
| 410H             | 1040    | IA32_MC4_CTL   | MC4_CTL             | If IA32_MCG_CAP.CNT >4  |
| 411H             | 1041    | IA32_MC4_STATUS  | MC4_STATUS          | If IA32_MCG_CAP.CNT >4  |
| 412H             | 1042    | IA32_MC4_ADDR <sup>1</sup>                               | MC4_ADDR            | If IA32_MCG_CAP.CNT >4  |
| 413H             | 1043    | IA32_MC4_MISC  | MC4_MISC            | If IA32_MCG_CAP.CNT >4  |
| 414H             | 1044    | IA32_MC5_CTL   | MC5_CTL             | If IA32_MCG_CAP.CNT >5  |
| 415H             | 1045    | IA32_MC5_STATUS  | MC5_STATUS          | If IA32_MCG_CAP.CNT >5  |
| 416H             | 1046    | IA32_MC5_ADDR <sup>1</sup>                               | MC5_ADDR            | If IA32_MCG_CAP.CNT >5  |
| 417H             | 1047    | IA32_MC5_MISC  | MC5_MISC            | If IA32_MCG_CAP.CNT >5  |
| 418H             | 1048    | IA32_MC6_CTL   | MC6_CTL             | If IA32_MCG_CAP.CNT >6  |
| 419H             | 1049    | IA32_MC6_STATUS  | MC6_STATUS          | If IA32_MCG_CAP.CNT >6  |
| 41AH             | 1050    | IA32_MC6_ADDR <sup>1</sup>                               | MC6_ADDR            | If IA32_MCG_CAP.CNT >6  |
| 41BH             | 1051    | IA32_MC6_MISC  | MC6_MISC            | If IA32_MCG_CAP.CNT >6  |
| 41CH             | 1052    | IA32_MC7_CTL   | MC7_CTL             | If IA32_MCG_CAP.CNT >7  |
| 41DH             | 1053    | IA32_MC7_STATUS  | MC7_STATUS          | If IA32_MCG_CAP.CNT >7  |
| 41EH             | 1054    | IA32_MC7_ADDR <sup>1</sup>                               | MC7_ADDR            | If IA32_MCG_CAP.CNT >7  |
| 41FH             | 1055    | IA32_MC7_MISC  | MC7_MISC            | If IA32_MCG_CAP.CNT >7  |
| 420H             | 1056    | IA32_MC8_CTL   | MC8_CTL             | If IA32_MCG_CAP.CNT >8  |
| 421H             | 1057    | IA32_MC8_STATUS  | MC8_STATUS          | If IA32_MCG_CAP.CNT >8  |
| 422H             | 1058    | IA32_MC8_ADDR <sup>1</sup>                               | MC8_ADDR            | If IA32_MCG_CAP.CNT >8  |
| 423H             | 1059    | IA32_MC8_MISC  | MC8_MISC            | If IA32_MCG_CAP.CNT >8  |
| 424H             | 1060    | IA32_MC9_CTL   | MC9_CTL             | If IA32_MCG_CAP.CNT >9  |
| 425H             | 1061    | IA32_MC9_STATUS  | MC9_STATUS          | If IA32_MCG_CAP.CNT >9  |
| 426H             | 1062    | IA32_MC9_ADDR <sup>1</sup>                               | MC9_ADDR            | If IA32_MCG_CAP.CNT >9  |
| 427H             | 1063    | IA32_MC9_MISC  | MC9_MISC            | If IA32_MCG_CAP.CNT >9  |
| 428H             | 1064    | IA32_MC10_CTL  | MC10_CTL            | If IA32_MCG_CAP.CNT >10 |
| 429H             | 1065    | IA32_MC10_STATUS   | MC10_STATUS         | If IA32_MCG_CAP.CNT >10 |
| 42AH             | 1066    | IA32_MC10_ADDR <sup>1</sup>                              | MC10_ADDR           | If IA32_MCG_CAP.CNT >10 |
| 42BH             | 1067    | IA32_MC10_MISC   | MC10_MISC           | If IA32_MCG_CAP.CNT >10 |
| 42CH             | 1068    | IA32_MC11_CTL  | MC11_CTL            | If IA32_MCG_CAP.CNT >11 |
| 42DH             | 1069    | IA32_MC11_STATUS   | MC11_STATUS         | If IA32_MCG_CAP.CNT >11 |
| 42EH             | 1070    | IA32_MC11_ADDR <sup>1</sup>                              | MC11_ADDR           | If IA32_MCG_CAP.CNT >11 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description | Comment                 |
|------------------|---------|--|---------------------|-------------------------|
| Hex              | Decimal |  |                     |                         |
| 42FH             | 1071    | IA32_MC11_MISC   | MC11_MISC           | If IA32_MCG_CAP.CNT >11 |
| 430H             | 1072    | IA32_MC12_CTL  | MC12_CTL            | If IA32_MCG_CAP.CNT >12 |
| 431H             | 1073    | IA32_MC12_STATUS   | MC12_STATUS         | If IA32_MCG_CAP.CNT >12 |
| 432H             | 1074    | IA32_MC12_ADDR <sup>7</sup>                              | MC12_ADDR           | If IA32_MCG_CAP.CNT >12 |
| 433H             | 1075    | IA32_MC12_MISC   | MC12_MISC           | If IA32_MCG_CAP.CNT >12 |
| 434H             | 1076    | IA32_MC13_CTL  | MC13_CTL            | If IA32_MCG_CAP.CNT >13 |
| 435H             | 1077    | IA32_MC13_STATUS   | MC13_STATUS         | If IA32_MCG_CAP.CNT >13 |
| 436H             | 1078    | IA32_MC13_ADDR <sup>7</sup>                              | MC13_ADDR           | If IA32_MCG_CAP.CNT >13 |
| 437H             | 1079    | IA32_MC13_MISC   | MC13_MISC           | If IA32_MCG_CAP.CNT >13 |
| 438H             | 1080    | IA32_MC14_CTL  | MC14_CTL            | If IA32_MCG_CAP.CNT >14 |
| 439H             | 1081    | IA32_MC14_STATUS   | MC14_STATUS         | If IA32_MCG_CAP.CNT >14 |
| 43AH             | 1082    | IA32_MC14_ADDR <sup>7</sup>                              | MC14_ADDR           | If IA32_MCG_CAP.CNT >14 |
| 43BH             | 1083    | IA32_MC14_MISC   | MC14_MISC           | If IA32_MCG_CAP.CNT >14 |
| 43CH             | 1084    | IA32_MC15_CTL  | MC15_CTL            | If IA32_MCG_CAP.CNT >15 |
| 43DH             | 1085    | IA32_MC15_STATUS   | MC15_STATUS         | If IA32_MCG_CAP.CNT >15 |
| 43EH             | 1086    | IA32_MC15_ADDR <sup>7</sup>                              | MC15_ADDR           | If IA32_MCG_CAP.CNT >15 |
| 43FH             | 1087    | IA32_MC15_MISC   | MC15_MISC           | If IA32_MCG_CAP.CNT >15 |
| 440H             | 1088    | IA32_MC16_CTL  | MC16_CTL            | If IA32_MCG_CAP.CNT >16 |
| 441H             | 1089    | IA32_MC16_STATUS   | MC16_STATUS         | If IA32_MCG_CAP.CNT >16 |
| 442H             | 1090    | IA32_MC16_ADDR <sup>7</sup>                              | MC16_ADDR           | If IA32_MCG_CAP.CNT >16 |
| 443H             | 1091    | IA32_MC16_MISC   | MC16_MISC           | If IA32_MCG_CAP.CNT >16 |
| 444H             | 1092    | IA32_MC17_CTL  | MC17_CTL            | If IA32_MCG_CAP.CNT >17 |
| 445H             | 1093    | IA32_MC17_STATUS   | MC17_STATUS         | If IA32_MCG_CAP.CNT >17 |
| 446H             | 1094    | IA32_MC17_ADDR <sup>7</sup>                              | MC17_ADDR           | If IA32_MCG_CAP.CNT >17 |
| 447H             | 1095    | IA32_MC17_MISC   | MC17_MISC           | If IA32_MCG_CAP.CNT >17 |
| 448H             | 1096    | IA32_MC18_CTL  | MC18_CTL            | If IA32_MCG_CAP.CNT >18 |
| 449H             | 1097    | IA32_MC18_STATUS   | MC18_STATUS         | If IA32_MCG_CAP.CNT >18 |
| 44AH             | 1098    | IA32_MC18_ADDR <sup>7</sup>                              | MC18_ADDR           | If IA32_MCG_CAP.CNT >18 |
| 44BH             | 1099    | IA32_MC18_MISC   | MC18_MISC           | If IA32_MCG_CAP.CNT >18 |
| 44CH             | 1100    | IA32_MC19_CTL  | MC19_CTL            | If IA32_MCG_CAP.CNT >19 |
| 44DH             | 1101    | IA32_MC19_STATUS   | MC19_STATUS         | If IA32_MCG_CAP.CNT >19 |
| 44EH             | 1102    | IA32_MC19_ADDR <sup>7</sup>                              | MC19_ADDR           | If IA32_MCG_CAP.CNT >19 |
| 44FH             | 1103    | IA32_MC19_MISC   | MC19_MISC           | If IA32_MCG_CAP.CNT >19 |
| 450H             | 1104    | IA32_MC20_CTL  | MC20_CTL            | If IA32_MCG_CAP.CNT >20 |
| 451H             | 1105    | IA32_MC20_STATUS   | MC20_STATUS         | If IA32_MCG_CAP.CNT >20 |
| 452H             | 1106    | IA32_MC20_ADDR <sup>7</sup>                              | MC20_ADDR           | If IA32_MCG_CAP.CNT >20 |
| 453H             | 1107    | IA32_MC20_MISC   | MC20_MISC           | If IA32_MCG_CAP.CNT >20 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment                  |
|------------------|---------|--|--|--------------------------|
| Hex              | Decimal |  |  |                          |
| 454H             | 1108    | IA32_MC21_CTL  | MC21_CTL   | If IA32_MCG_CAP.CNT >21  |
| 455H             | 1109    | IA32_MC21_STATUS   | MC21_STATUS  | If IA32_MCG_CAP.CNT >21  |
| 456H             | 1110    | IA32_MC21_ADDR <sup>1</sup>                              | MC21_ADDR  | If IA32_MCG_CAP.CNT >21  |
| 457H             | 1111    | IA32_MC21_MISC   | MC21_MISC  | If IA32_MCG_CAP.CNT >21  |
| 458H             | 1112    | IA32_MC22_CTL  | MC22_CTL   | If IA32_MCG_CAP.CNT >22  |
| 459H             | 1113    | IA32_MC22_STATUS   | MC22_STATUS  | If IA32_MCG_CAP.CNT >22  |
| 45AH             | 1114    | IA32_MC22_ADDR <sup>1</sup>                              | MC22_ADDR  | If IA32_MCG_CAP.CNT >22  |
| 45BH             | 1115    | IA32_MC22_MISC   | MC22_MISC  | If IA32_MCG_CAP.CNT >22  |
| 45CH             | 1116    | IA32_MC23_CTL  | MC23_CTL   | If IA32_MCG_CAP.CNT >23  |
| 45DH             | 1117    | IA32_MC23_STATUS   | MC23_STATUS  | If IA32_MCG_CAP.CNT >23  |
| 45EH             | 1118    | IA32_MC23_ADDR <sup>1</sup>                              | MC23_ADDR  | If IA32_MCG_CAP.CNT >23  |
| 45FH             | 1119    | IA32_MC23_MISC   | MC23_MISC  | If IA32_MCG_CAP.CNT >23  |
| 460H             | 1120    | IA32_MC24_CTL  | MC24_CTL   | If IA32_MCG_CAP.CNT >24  |
| 461H             | 1121    | IA32_MC24_STATUS   | MC24_STATUS  | If IA32_MCG_CAP.CNT >24  |
| 462H             | 1122    | IA32_MC24_ADDR <sup>1</sup>                              | MC24_ADDR  | If IA32_MCG_CAP.CNT >24  |
| 463H             | 1123    | IA32_MC24_MISC   | MC24_MISC  | If IA32_MCG_CAP.CNT >24  |
| 464H             | 1124    | IA32_MC25_CTL  | MC25_CTL   | If IA32_MCG_CAP.CNT >25  |
| 465H             | 1125    | IA32_MC25_STATUS   | MC25_STATUS  | If IA32_MCG_CAP.CNT >25  |
| 466H             | 1126    | IA32_MC25_ADDR <sup>1</sup>                              | MC25_ADDR  | If IA32_MCG_CAP.CNT >25  |
| 467H             | 1127    | IA32_MC25_MISC   | MC25_MISC  | If IA32_MCG_CAP.CNT >25  |
| 468H             | 1128    | IA32_MC26_CTL  | MC26_CTL   | If IA32_MCG_CAP.CNT >26  |
| 469H             | 1129    | IA32_MC26_STATUS   | MC26_STATUS  | If IA32_MCG_CAP.CNT >26  |
| 46AH             | 1130    | IA32_MC26_ADDR <sup>1</sup>                              | MC26_ADDR  | If IA32_MCG_CAP.CNT >26  |
| 46BH             | 1131    | IA32_MC26_MISC   | MC26_MISC  | If IA32_MCG_CAP.CNT >26  |
| 46CH             | 1132    | IA32_MC27_CTL  | MC27_CTL   | If IA32_MCG_CAP.CNT >27  |
| 46DH             | 1133    | IA32_MC27_STATUS   | MC27_STATUS  | If IA32_MCG_CAP.CNT >27  |
| 46EH             | 1134    | IA32_MC27_ADDR <sup>1</sup>                              | MC27_ADDR  | If IA32_MCG_CAP.CNT >27  |
| 46FH             | 1135    | IA32_MC27_MISC   | MC27_MISC  | If IA32_MCG_CAP.CNT >27  |
| 470H             | 1136    | IA32_MC28_CTL  | MC28_CTL   | If IA32_MCG_CAP.CNT >28  |
| 471H             | 1137    | IA32_MC28_STATUS   | MC28_STATUS  | If IA32_MCG_CAP.CNT >28  |
| 472H             | 1138    | IA32_MC28_ADDR <sup>1</sup>                              | MC28_ADDR  | If IA32_MCG_CAP.CNT >28  |
| 473H             | 1139    | IA32_MC28_MISC   | MC28_MISC  | If IA32_MCG_CAP.CNT >28  |
| 480H             | 1152    | IA32_VMX_BASIC   | Reporting Register of Basic VMX Capabilities (R/O)<br>See Appendix A.1, "Basic VMX Information." | If CPUID.01H:ECX.[5] = 1 |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment  |
|------------------|---------|--|--|--|
| Hex              | Decimal |  |  |  |
| 481H             | 1153    | IA32_VMX_PINBASED_CTL5                                   | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Appendix A.3.1, "Pin-Based VM-Execution Controls."                                 | If CPUID.01H:ECX.[5] = 1   |
| 482H             | 1154    | IA32_VMX_PROCBASED_CTL5                                  | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3.2, "Primary Processor-Based VM-Execution Controls."     | If CPUID.01H:ECX.[5] = 1   |
| 483H             | 1155    | IA32_VMX_EXIT_CTL5                                       | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Appendix A.4, "VM-Exit Controls."   | If CPUID.01H:ECX.[5] = 1   |
| 484H             | 1156    | IA32_VMX_ENTRY_CTL5                                      | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Appendix A.5, "VM-Entry Controls."   | If CPUID.01H:ECX.[5] = 1   |
| 485H             | 1157    | IA32_VMX_MISC  | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Appendix A.6, "Miscellaneous Data."  | If CPUID.01H:ECX.[5] = 1   |
| 486H             | 1158    | IA32_VMX_CRO_FIXED0                                      | Capability Reporting Register of CRO Bits Fixed to 0 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CRO."   | If CPUID.01H:ECX.[5] = 1   |
| 487H             | 1159    | IA32_VMX_CRO_FIXED1                                      | Capability Reporting Register of CRO Bits Fixed to 1 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CRO."   | If CPUID.01H:ECX.[5] = 1   |
| 488H             | 1160    | IA32_VMX_CR4_FIXED0                                      | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4."   | If CPUID.01H:ECX.[5] = 1   |
| 489H             | 1161    | IA32_VMX_CR4_FIXED1                                      | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4."   | If CPUID.01H:ECX.[5] = 1   |
| 48AH             | 1162    | IA32_VMX_VMCS_ENUM                                       | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Appendix A.9, "VMCS Enumeration."   | If CPUID.01H:ECX.[5] = 1   |
| 48BH             | 1163    | IA32_VMX_PROCBASED_CTL52                                 | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3.3, "Secondary Processor-Based VM-Execution Controls." | If ( CPUID.01H:ECX.[5] && IA32_VMX_PROCBASED_CTL5[63])   |
| 48CH             | 1164    | IA32_VMX_EPT_VPID_CAP                                    | Capability Reporting Register of EPT and VPID (R/O)<br>See Appendix A.10, "VPID and EPT Capabilities."   | If ( CPUID.01H:ECX.[5] && IA32_VMX_PROCBASED_CTL5[63] && ( IA32_VMX_PROCBASED_CTL52[33]    IA32_VMX_PROCBASED_CTL52[37]) ) |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment  |
|------------------|---------|--|---|--|
| Hex              | Decimal |  |   |  |
| 48DH             | 1165    | IA32_VMX_TRUE_PINBASED_CTL                               | Capability Reporting Register of Pin-Based VM-Execution Flex Controls (R/O)<br>See Appendix A.3.1, "Pin-Based VM-Execution Controls."                             | If ( CPUID.01H:ECX.[5] = 1 && IA32_VMX_BASIC[55] )                 |
| 48EH             | 1166    | IA32_VMX_TRUE_PROCBASED_CTL                              | Capability Reporting Register of Primary Processor-Based VM-Execution Flex Controls (R/O)<br>See Appendix A.3.2, "Primary Processor-Based VM-Execution Controls." | If ( CPUID.01H:ECX.[5] = 1 && IA32_VMX_BASIC[55] )                 |
| 48FH             | 1167    | IA32_VMX_TRUE_EXIT_CTL                                   | Capability Reporting Register of VM-Exit Flex Controls (R/O)<br>See Appendix A.4, "VM-Exit Controls."   | If ( CPUID.01H:ECX.[5] = 1 && IA32_VMX_BASIC[55] )                 |
| 490H             | 1168    | IA32_VMX_TRUE_ENTRY_CTL                                  | Capability Reporting Register of VM-Entry Flex Controls (R/O)<br>See Appendix A.5, "VM-Entry Controls."   | If ( CPUID.01H:ECX.[5] = 1 && IA32_VMX_BASIC[55] )                 |
| 491H             | 1169    | IA32_VMX_VMFUNC  | Capability Reporting Register of VM-Function Controls (R/O)   | If ( CPUID.01H:ECX.[5] = 1 && IA32_VMX_BASIC[55] )                 |
| 4C1H             | 1217    | IA32_A_PMC0  | Full Width Writable IA32_PMC0 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 0) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |
| 4C2H             | 1218    | IA32_A_PMC1  | Full Width Writable IA32_PMC1 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 1) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |
| 4C3H             | 1219    | IA32_A_PMC2  | Full Width Writable IA32_PMC2 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 2) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |
| 4C4H             | 1220    | IA32_A_PMC3  | Full Width Writable IA32_PMC3 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 3) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |
| 4C5H             | 1221    | IA32_A_PMC4  | Full Width Writable IA32_PMC4 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 4) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |
| 4C6H             | 1222    | IA32_A_PMC5  | Full Width Writable IA32_PMC5 Alias (R/W)   | (If CPUID.0AH: EAX[15:8] > 5) &&<br>IA32_PERF_CAPABILITIES[13] = 1 |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
| 4C7H             | 1223    | IA32_A_PMC6  | Full Width Writable IA32_PMC6 Alias (R/W)  | (If CPUID.0AH: EAX[15:8] > 6) &&<br>IA32_PERF_CAPABILITIES[13] = 1  |
| 4C8H             | 1224    | IA32_A_PMC7  | Full Width Writable IA32_PMC7 Alias (R/W)  | (If CPUID.0AH: EAX[15:8] > 7) &&<br>IA32_PERF_CAPABILITIES[13] = 1  |
| 4D0H             | 1232    | IA32_MCG_EXT_CTL   | Allows software to signal some MCEs to only a single logical processor in the system. (R/W)<br>See Section 15.3.1.4, "IA32_MCG_EXT_CTL MSR". | If IA32_MCG_CAP.LMCE_P = 1  |
|                  |         | 0  | LMCE_EN  |   |
|                  |         | 63:1   | Reserved   |   |
| 500H             | 1280    | IA32_SGX_SVN_STATUS                                      | Status and SVN Threshold of SGX Support for ACM (RO).  | If CPUID.(EAX=07H, ECX=0H): EBX[2] = 1  |
|                  |         | 0  | Lock   | See Section 41.11.3, "Interactions with Authenticated Code Modules (ACMs)".   |
|                  |         | 15:1   | Reserved   |   |
|                  |         | 23:16  | SGX_SVN_SINIT  | See Section 41.11.3, "Interactions with Authenticated Code Modules (ACMs)".   |
|                  |         | 63:24  | Reserved   |   |
| 560H             | 1376    | IA32_RTIT_OUTPUT_BASE                                    | Trace Output Base Register (R/W)   | If ((CPUID.(EAX=07H, ECX=0):EBX[25] = 1) && (CPUID.(EAX=14H,ECX=0):ECX[0] = 1)    (CPUID.(EAX=14H,ECX=0):ECX[2] = 1)) |
|                  |         | 6:0  | Reserved   |   |
|                  |         | MAXPHYADDR <sup>3</sup> -1:7                             | Base physical address.   |   |
|                  |         | 63:MAXPHYADDR  | Reserved   |   |
| 561H             | 1377    | IA32_RTIT_OUTPUT_MASK_PTRS                               | Trace Output Mask Pointers Register (R/W)  | If ((CPUID.(EAX=07H, ECX=0):EBX[25] = 1) && (CPUID.(EAX=14H,ECX=0):ECX[0] = 1)    (CPUID.(EAX=14H,ECX=0):ECX[2] = 1)) |
|                  |         | 6:0  | Reserved   |   |
|                  |         | 31:7   | MaskOrTableOffset  |   |
|                  |         | 63:32  | Output Offset  |   |



Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |                         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description           | Comment                                  |
|------------------|-------------------------|--|-------------------------------|--|
| Hex              | Decimal                 |  |                               |  |
| 570H             | 1392                    | IA32_RTIT_CTL  | Trace Control Register (R/W)  | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1)  |
|                  |                         | 0  | TraceEn                       |  |
|                  |                         | 1  | CYCEn                         | If (CPUID.(EAX=07H, ECX=0):EBX[1] = 1)   |
|                  |                         | 2  | OS                            |  |
|                  |                         | 3  | User                          |  |
|                  |                         | 4  | PwrEvtEn                      |  |
|                  |                         | 5  | FUPonPTW                      |  |
|                  |                         | 6  | FabricEn                      | If (CPUID.(EAX=07H, ECX=0):ECX[3] = 1)   |
|                  |                         | 7  | CR3 filter                    |  |
|                  |                         | 8  | ToPA                          |  |
|                  |                         | 9  | MTCEn                         | If (CPUID.(EAX=07H, ECX=0):EBX[3] = 1)   |
|                  |                         | 10   | TSCEn                         |  |
|                  |                         | 11   | DisRETC                       |  |
|                  |                         | 12   | PTWEn                         |  |
|                  |                         | 13   | BranchEn                      |  |
|                  |                         | 17:14  | MTCFreq                       | If (CPUID.(EAX=07H, ECX=0):EBX[3] = 1)   |
|                  |                         | 18   | Reserved, must be zero.       |  |
|                  |                         | 22:19  | CYCThresh                     | If (CPUID.(EAX=07H, ECX=0):EBX[1] = 1)   |
|                  |                         | 23   | Reserved, must be zero.       |  |
|                  |                         | 27:24  | PSBFreq                       | If (CPUID.(EAX=07H, ECX=0):EBX[1] = 1)   |
|                  |                         | 31:28  | Reserved, must be zero.       |  |
|                  |                         | 35:32  | ADDR0_CFG                     | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 0) |
|                  |                         | 39:36  | ADDR1_CFG                     | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 1) |
| 43:40            | ADDR2_CFG               | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 2)                 |                               |  |
| 47:44            | ADDR3_CFG               | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 3)                 |                               |  |
| 63:48            | Reserved, must be zero. |  |                               |  |
| 571H             | 1393                    | IA32_RTIT_STATUS   | Tracing Status Register (R/W) | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1)  |
|                  |                         | 0  | FilterEn (writes ignored)     | If (CPUID.(EAX=07H, ECX=0):EBX[2] = 1)   |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                   | Comment                                  |
|------------------|---------|--|---------------------------------------|--|
| Hex              | Decimal |  |                                       |  |
|                  |         | 1  | ContexEn (writes ignored)             |  |
|                  |         | 2  | TriggerEn (writes ignored)            |  |
|                  |         | 3  | Reserved                              |  |
|                  |         | 4  | Error                                 |  |
|                  |         | 5  | Stopped                               |  |
|                  |         | 31:6   | Reserved, must be zero.               |  |
|                  |         | 48:32  | PacketByteCnt                         | If (CPUID.(EAX=07H, ECX=0):EBX[1] > 3)   |
|                  |         | 63:49  | Reserved                              |  |
| 572H             | 1394    | IA32_RTIT_CR3_MATCH                                      | Trace Filter CR3 Match Register (R/W) | If (CPUID.(EAX=07H, ECX=0):EBX[25] = 1)  |
|                  |         | 4:0  | Reserved                              |  |
|                  |         | 63:5   | CR3[63:5] value to match.             |  |
| 580H             | 1408    | IA32_RTIT_ADDR0_A  | Region 0 Start Address (R/W)          | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 0) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 581H             | 1409    | IA32_RTIT_ADDR0_B  | Region 0 End Address (R/W)            | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 0) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 582H             | 1410    | IA32_RTIT_ADDR1_A  | Region 1 Start Address (R/W)          | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 1) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 583H             | 1411    | IA32_RTIT_ADDR1_B  | Region 1 End Address (R/W)            | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 1) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 584H             | 1412    | IA32_RTIT_ADDR2_A  | Region 2 Start Address (R/W)          | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 2) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 585H             | 1413    | IA32_RTIT_ADDR2_B  | Region 2 End Address (R/W)            | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 2) |
|                  |         | 47:0   | Virtual Address                       |  |
|                  |         | 63:48  | SignExt_VA                            |  |
| 586H             | 1414    | IA32_RTIT_ADDR3_A  | Region 3 Start Address (R/W)          | If (CPUID.(EAX=07H, ECX=1):EAX[2:0] > 3) |
|                  |         | 47:0   | Virtual Address                       |  |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment                                  |
|------------------|---------|--|---|--|
| Hex              | Decimal |  |   |  |
|                  |         | 63:48  | SignExt_VA  |  |
| 587H             | 1415    | IA32_RTIT_ADDR3_B  | Region 3 End Address (R/W)  | If (CPUID.(EAX=07H, ECX=1);EAX[2:0] > 3) |
|                  |         | 47:0   | Virtual Address   |  |
|                  |         | 63:48  | SignExt_VA  |  |
| 600H             | 1536    | IA32_DS_AREA   | DS Save Area (R/W)<br>Points to the linear address of the first byte of the DS buffer management area, which is used to manage the BTS and PEBS buffers.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism." | If (CPUID.01H:EDX.DS[21] = 1)            |
|                  |         | 63:0   | The linear address of the first byte of the DS buffer management area, if IA-32e mode is active.  |  |
|                  |         | 31:0   | The linear address of the first byte of the DS buffer management area, if not in IA-32e mode.   |  |
|                  |         | 63:32  | Reserved if not in IA-32e mode.   |  |
| 6E0H             | 1760    | IA32_TSC_DEADLINE  | TSC Target of Local APIC's TSC Deadline Mode (R/W)  | If CPUID.01H:ECX.[24] = 1                |
| 770H             | 1904    | IA32_PM_ENABLE   | Enable/disable HWP (R/W)  | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 0  | HWP_ENABLE (R/W1-Once)<br>See Section 14.4.2, "Enabling HWP".   | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 63:1   | Reserved  |  |
| 771H             | 1905    | IA32_HWP_CAPABILITIES                                    | HWP Performance Range Enumeration (RO)  | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 7:0  | Highest_Performance<br>See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".  | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 15:8   | Guaranteed_Performance<br>See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".   | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 23:16  | Most_Efficient_Performance<br>See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".   | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 31:24  | Lowest_Performance<br>See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".   | If CPUID.06H:EAX.[7] = 1                 |
|                  |         | 63:32  | Reserved  |  |
| 772H             | 1906    | IA32_HWP_REQUEST_PKG                                     | Power Management Control Hints for All Logical Processors in a Package (R/W)  | If CPUID.06H:EAX.[11] = 1                |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
|                  |         | 7:0  | Minimum_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[11] = 1                                 |
|                  |         | 15:8   | Maximum_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[11] = 1                                 |
|                  |         | 23:16  | Desired_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[11] = 1                                 |
|                  |         | 31:24  | Energy_Performance_Preference<br>See Section 14.4.4, "Managing HWP".         | If CPUID.06H:EAX.[11] = 1<br>&&<br>CPUID.06H:EAX.[10] = 1 |
|                  |         | 41:32  | Activity_Window<br>See Section 14.4.4, "Managing HWP".                       | If CPUID.06H:EAX.[11] = 1<br>&&<br>CPUID.06H:EAX.[9] = 1  |
|                  |         | 63:42  | Reserved   |   |
| 773H             | 1907    | IA32_HWP_INTERRUPT                                       | Control HWP Native Interrupts (R/W)  | If CPUID.06H:EAX.[8] = 1                                  |
|                  |         | 0  | EN_Guaranteed_Performance_Change<br>See Section 14.4.6, "HWP Notifications". | If CPUID.06H:EAX.[8] = 1                                  |
|                  |         | 1  | EN_Excursion_Minimum<br>See Section 14.4.6, "HWP Notifications".             | If CPUID.06H:EAX.[8] = 1                                  |
|                  |         | 63:2   | Reserved   |   |
| 774H             | 1908    | IA32_HWP_REQUEST   | Power Management Control Hints to a Logical Processor (R/W)                  | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 7:0  | Minimum_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 15:8   | Maximum_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 23:16  | Desired_Performance<br>See Section 14.4.4, "Managing HWP".                   | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 31:24  | Energy_Performance_Preference<br>See Section 14.4.4, "Managing HWP".         | If CPUID.06H:EAX.[7] = 1<br>&& CPUID.06H:EAX.[10] = 1     |
|                  |         | 41:32  | Activity_Window<br>See Section 14.4.4, "Managing HWP".                       | If CPUID.06H:EAX.[7] = 1<br>&& CPUID.06H:EAX.[9] = 1      |
|                  |         | 42   | Package_Control<br>See Section 14.4.4, "Managing HWP".                       | If CPUID.06H:EAX.[7] = 1<br>&& CPUID.06H:EAX.[11] = 1     |
|                  |         | 63:43  | Reserved   |   |
| 777H             | 1911    | IA32_HWP_STATUS  | Log bits indicating changes to Guaranteed & excursions to Minimum (R/W)      | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 0  | Guaranteed_Performance_Change (R/WCO)<br>See Section 14.4.5, "HWP Feedback". | If CPUID.06H:EAX.[7] = 1                                  |
|                  |         | 1  | Reserved   |   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment   |
|------------------|---------|--|---|---|
| Hex              | Decimal |  |   |   |
|                  |         | 2  | Excursion_To_Minimum (R/WCO)<br>See Section 14.4.5, "HWP Feedback". | If CPUID.06H:EAX.[7] = 1                                |
|                  |         | 63:3   | Reserved  |   |
| 802H             | 2050    | IA32_X2APIC_APICID                                       | x2APIC ID Register (R/O)<br>See x2APIC Specification.               | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 803H             | 2051    | IA32_X2APIC_VERSION                                      | x2APIC Version Register (R/O)                                       | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 808H             | 2056    | IA32_X2APIC_TPR  | x2APIC Task Priority Register (R/W)                                 | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 80AH             | 2058    | IA32_X2APIC_PPR  | x2APIC Processor Priority Register (R/O)                            | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 80BH             | 2059    | IA32_X2APIC_EOI  | x2APIC EOI Register (W/O)   | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 80DH             | 2061    | IA32_X2APIC_LDR  | x2APIC Logical Destination Register (R/O)                           | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 80FH             | 2063    | IA32_X2APIC_SIVR   | x2APIC Spurious Interrupt Vector Register (R/W)                     | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 810H             | 2064    | IA32_X2APIC_ISR0   | x2APIC In-Service Register Bits 31:0 (R/O)                          | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 811H             | 2065    | IA32_X2APIC_ISR1   | x2APIC In-Service Register Bits 63:32 (R/O)                         | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 812H             | 2066    | IA32_X2APIC_ISR2   | x2APIC In-Service Register Bits 95:64 (R/O)                         | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 813H             | 2067    | IA32_X2APIC_ISR3   | x2APIC In-Service Register Bits 127:96 (R/O)                        | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 814H             | 2068    | IA32_X2APIC_ISR4   | x2APIC In-Service Register Bits 159:128 (R/O)                       | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 815H             | 2069    | IA32_X2APIC_ISR5   | x2APIC In-Service Register Bits 191:160 (R/O)                       | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |
| 816H             | 2070    | IA32_X2APIC_ISR6   | x2APIC In-Service Register Bits 223:192 (R/O)                       | If CPUID.01H:ECX.[21] = 1<br>&& IA32_APIC_BASE.[10] = 1 |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description                                  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
| 817H             | 2071    | IA32_X2APIC_ISR7   | x2APIC In-Service Register Bits 255:224 (R/O)        | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 818H             | 2072    | IA32_X2APIC_TMR0   | x2APIC Trigger Mode Register Bits 31:0 (R/O)         | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 819H             | 2073    | IA32_X2APIC_TMR1   | x2APIC Trigger Mode Register Bits 63:32 (R/O)        | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 81AH             | 2074    | IA32_X2APIC_TMR2   | x2APIC Trigger Mode Register Bits 95:64 (R/O)        | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 81BH             | 2075    | IA32_X2APIC_TMR3   | x2APIC Trigger Mode Register Bits 127:96 (R/O)       | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 81CH             | 2076    | IA32_X2APIC_TMR4   | x2APIC Trigger Mode Register Bits 159:128 (R/O)      | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 81DH             | 2077    | IA32_X2APIC_TMR5   | x2APIC Trigger Mode Register Bits 191:160 (R/O)      | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 81EH             | 2078    | IA32_X2APIC_TMR6   | x2APIC Trigger Mode Register Bits 223:192 (R/O)      | If ( CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1) |
| 81FH             | 2079    | IA32_X2APIC_TMR7   | x2APIC Trigger Mode Register Bits 255:224 (R/O)      | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 820H             | 2080    | IA32_X2APIC_IRR0   | x2APIC Interrupt Request Register Bits 31:0 (R/O)    | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 821H             | 2081    | IA32_X2APIC_IRR1   | x2APIC Interrupt Request Register Bits 63:32 (R/O)   | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 822H             | 2082    | IA32_X2APIC_IRR2   | x2APIC Interrupt Request Register Bits 95:64 (R/O)   | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 823H             | 2083    | IA32_X2APIC_IRR3   | x2APIC Interrupt Request Register Bits 127:96 (R/O)  | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 824H             | 2084    | IA32_X2APIC_IRR4   | x2APIC Interrupt Request Register Bits 159:128 (R/O) | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |
| 825H             | 2085    | IA32_X2APIC_IRR5   | x2APIC Interrupt Request Register Bits 191:160 (R/O) | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1    |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment  |
|------------------|---------|--|---|--|
| Hex              | Decimal |  |   |  |
| 826H             | 2086    | IA32_X2APIC_IRR6   | x2APIC Interrupt Request Register Bits 223:192 (R/O)        | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 827H             | 2087    | IA32_X2APIC_IRR7   | x2APIC Interrupt Request Register Bits 255:224 (R/O)        | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 828H             | 2088    | IA32_X2APIC_ESR  | x2APIC Error Status Register (R/w)                          | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 82FH             | 2095    | IA32_X2APIC_LVT_CMCI                                     | x2APIC LVT Corrected Machine Check Interrupt Register (R/w) | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 830H             | 2096    | IA32_X2APIC_ICR  | x2APIC Interrupt Command Register (R/w)                     | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 832H             | 2098    | IA32_X2APIC_LVT_TIMER                                    | x2APIC LVT Timer Interrupt Register (R/w)                   | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 833H             | 2099    | IA32_X2APIC_LVT_THERMAL                                  | x2APIC LVT Thermal Sensor Interrupt Register (R/w)          | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 834H             | 2100    | IA32_X2APIC_LVT_PMI                                      | x2APIC LVT Performance Monitor Interrupt Register (R/w)     | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 835H             | 2101    | IA32_X2APIC_LVT_LINT0                                    | x2APIC LVT LINT0 Register (R/w)                             | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 836H             | 2102    | IA32_X2APIC_LVT_LINT1                                    | x2APIC LVT LINT1 Register (R/w)                             | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 837H             | 2103    | IA32_X2APIC_LVT_ERROR                                    | x2APIC LVT Error Register (R/w)                             | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 838H             | 2104    | IA32_X2APIC_INIT_COUNT                                   | x2APIC Initial Count Register (R/w)                         | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 839H             | 2105    | IA32_X2APIC_CUR_COUNT                                    | x2APIC Current Count Register (R/O)                         | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 83EH             | 2110    | IA32_X2APIC_DIV_CONF                                     | x2APIC Divide Configuration Register (R/w)                  | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| 83FH             | 2111    | IA32_X2APIC_SELF_IPI                                     | x2APIC Self IPI Register (W/O)                              | If CPUID.01H:ECX.[21] = 1 && IA32_APIC_BASE.[10] = 1 |
| C80H             | 3200    | IA32_DEBUG_INTERFACE                                     | Silicon Debug Feature Control (R/w)                         | If CPUID.01H:ECX.[11] = 1                            |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment  |
|------------------|---------|--|---|--|
| Hex              | Decimal |  |   |  |
|                  |         | 0  | Enable (R/W)<br>BIOS set 1 to enable Silicon debug features.<br>Default is 0.   | If CPUID.01H:ECX.[11] = 1  |
|                  |         | 29:1   | Reserved  |  |
|                  |         | 30   | Lock (R/W): If 1, locks any further change to the MSR. The lock bit is set automatically on the first SMI assertion even if not explicitly set by BIOS. Default is 0. | If CPUID.01H:ECX.[11] = 1  |
|                  |         | 31   | Debug Occurred (R/O): This “sticky bit” is set by hardware to indicate the status of bit 0. Default is 0.   | If CPUID.01H:ECX.[11] = 1  |
|                  |         | 63:32  | Reserved  |  |
| C81H             | 3201    | IA32_L3_QOS_CFG  | L3 QOS Configuration (R/W)  | If ( CPUID.(EAX=10H, ECX=1):ECX.[2] = 1 )                            |
|                  |         | 0  | Enable (R/W)<br>Set 1 to enable L3 CAT masks and COS to operate in Code and Data Prioritization (CDP) mode.   |  |
|                  |         | 63:1   | Reserved. Attempts to write to reserved bits result in a #GP(0).  |  |
| C82H             | 3202    | IA32_L2_QOS_CFG  | L2 QOS Configuration (R/W)  | If ( CPUID.(EAX=10H, ECX=2):ECX.[2] = 1 )                            |
|                  |         | 0  | Enable (R/W)<br>Set 1 to enable L2 CAT masks and COS to operate in Code and Data Prioritization (CDP) mode.   |  |
|                  |         | 63:1   | Reserved. Attempts to write to reserved bits result in a #GP(0).  |  |
| C8DH             | 3213    | IA32_QM_EVTSEL   | Monitoring Event Select Register (R/W)  | If ( CPUID.(EAX=07H, ECX=0):EBX.[12] = 1 )                           |
|                  |         | 7:0  | Event ID: ID of a supported monitoring event to report via IA32_QM_CTR.   |  |
|                  |         | 31: 8  | Reserved  |  |
|                  |         | N+31:32  | Resource Monitoring ID: ID for monitoring hardware to report monitored data via IA32_QM_CTR.  | N = Ceil (Log <sub>2</sub> ( CPUID.(EAX= 0FH, ECX=0H).EBX[31:0] +1)) |
|                  |         | 63:N+32  | Reserved  |  |
| C8EH             | 3214    | IA32_QM_CTR  | Monitoring Counter Register (R/O)   | If ( CPUID.(EAX=07H, ECX=0):EBX.[12] = 1 )                           |
|                  |         | 61:0   | Resource Monitored Data   |  |
|                  |         | 62   | Unavailable: If 1, indicates data for this RMID is not available or not monitored for this resource or RMID.  |  |



Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|------------------|---------|--|--|---|
| Hex              | Decimal |  |  |   |
|                  |         | 63   | Error: If 1, indicates an unsupported RMID or event type was written to IA32_PQR_QM_EVTSEL.                |   |
| C8FH             | 3215    | IA32_PQR_ASSOC   | Resource Association Register (R/W)  | If ( (CPUID.(EAX=07H, ECX=0):EBX[12] = 1) or (CPUID.(EAX=07H, ECX=0):EBX[15] = 1) ) |
|                  |         | N-1:0  | Resource Monitoring ID (R/W): ID for monitoring hardware to track internal operation, e.g., memory access. | N = Ceil (Log <sub>2</sub> ( CPUID.(EAX= 0FH, ECX=0H).EBX[31:0] + 1))               |
|                  |         | 31:N   | Reserved   |   |
|                  |         | 63:32  | COS (R/W): The class of service (COS) to enforce (on writes); returns the current COS when read.           | If ( CPUID.(EAX=07H, ECX=0):EBX.[15] = 1 )  |
| C90H - D8FH      |         | Reserved MSR Address Space for CAT Mask Registers        | See Section 17.19.4.1, "Enumeration and Detection Support of Cache Allocation Technology".                 |   |
| C90H             | 3216    | IA32_L3_MASK_0   | L3 CAT Mask for COS0 (R/W)   | If (CPUID.(EAX=10H, ECX=0H):EBX[1] != 0)  |
|                  |         | 31:0   | Capacity Bit Mask (R/W)  |   |
|                  |         | 63:32  | Reserved   |   |
| C90H+n           | 3216+n  | IA32_L3_MASK_n   | L3 CAT Mask for COSn (R/W)   | n = CPUID.(EAX=10H, ECX=1H):EDX[15:0]   |
|                  |         | 31:0   | Capacity Bit Mask (R/W)  |   |
|                  |         | 63:32  | Reserved   |   |
| D10H - D4FH      |         | Reserved MSR Address Space for L2 CAT Mask Registers     | See Section 17.19.4.1, "Enumeration and Detection Support of Cache Allocation Technology".                 |   |
| D10H             | 3344    | IA32_L2_MASK_0   | L2 CAT Mask for COS0 (R/W)   | If (CPUID.(EAX=10H, ECX=0H):EBX[2] != 0)  |
|                  |         | 31:0   | Capacity Bit Mask (R/W)  |   |
|                  |         | 63:32  | Reserved   |   |
| D10H+n           | 3344+n  | IA32_L2_MASK_n   | L2 CAT Mask for COSn (R/W)   | n = CPUID.(EAX=10H, ECX=2H):EDX[15:0]   |
|                  |         | 31:0   | Capacity Bit Mask (R/W)  |   |
|                  |         | 63:32  | Reserved   |   |

**Table 2-2. IA-32 Architectural MSRs (Contd.)**

| Register Address        |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description  | Comment   |
|-------------------------|---------|--|--|---|
| Hex                     | Decimal |  |  |   |
| D90H                    | 3472    | IA32_BNDCFGS   | Supervisor State of MPX Configuration (R/W)  | If (CPUID.(EAX=07H, ECX=0H):EBX[14] = 1)                  |
|                         |         | 0  | EN: Enable Intel MPX in supervisor mode.   |   |
|                         |         | 1  | BNDPRESERVE: Preserve the bounds registers for near branch instructions in the absence of the BND prefix.  |   |
|                         |         | 11:2   | Reserved, must be zero.  |   |
|                         |         | 63:12  | Base Address of Bound Directory.   |   |
| DA0H                    | 3488    | IA32_XSS   | Extended Supervisor State Mask (R/W)   | If (CPUID.(0DH, 1):EAX.[3] = 1)                           |
|                         |         | 7:0  | Reserved   |   |
|                         |         | 8  | Trace Packet Configuration State (R/W)   |   |
|                         |         | 63:9   | Reserved.  |   |
| DB0H                    | 3504    | IA32_PKG_HDC_CTL   | Package Level Enable/disable HDC (R/W)   | If CPUID.06H:EAX.[13] = 1                                 |
|                         |         | 0  | HDC_Pkg_Enable (R/W)<br>Force HDC idling or wake up HDC-idled logical processors in the package. See Section 14.5.2, "Package level Enabling HDC". | If CPUID.06H:EAX.[13] = 1                                 |
|                         |         | 63:1   | Reserved   |   |
| DB1H                    | 3505    | IA32_PM_CTL1   | Enable/disable HWP (R/W)   | If CPUID.06H:EAX.[13] = 1                                 |
|                         |         | 0  | HDC_Allow_Block (R/W)<br>Allow/Block this logical processor for package level HDC control. See Section 14.5.3.                                     | If CPUID.06H:EAX.[13] = 1                                 |
|                         |         | 63:1   | Reserved   |   |
| DB2H                    | 3506    | IA32_THREAD_STALL  | Per-Logical_Processor HDC Idle Residency (R/O)   | If CPUID.06H:EAX.[13] = 1                                 |
|                         |         | 63:0   | Stall_Cycle_Cnt (R/W)<br>Stalled cycles due to HDC forced idle on this logical processor. See Section 14.5.4.1.                                    | If CPUID.06H:EAX.[13] = 1                                 |
| 4000_0000H - 4000_00FFH |         | Reserved MSR Address Space                               | All existing and future processors will not implement MSRs in this range.  |   |
| C000_0080H              |         | IA32_EFER  | Extended Feature Enables   | If (CPUID.80000001H:EDX.[20]    CPUID.80000001H:EDX.[29]) |
|                         |         | 0  | SYSCALL Enable: IA32_EFER.SCE (R/W)<br>Enables SYSCALL/SYSRET instructions in 64-bit mode.   |   |

Table 2-2. IA-32 Architectural MSRs (Contd.)

| Register Address |         | Architectural MSR Name / Bit Fields<br>(Former MSR Name) | MSR/Bit Description   | Comment   |
|------------------|---------|--|---|---|
| Hex              | Decimal |  |   |   |
|                  |         | 7:1  | Reserved  |   |
|                  |         | 8  | IA-32e Mode Enable: IA32_EFER.LME (R/W)<br>Enables IA-32e mode operation.   |   |
|                  |         | 9  | Reserved  |   |
|                  |         | 10   | IA-32e Mode Active: IA32_EFER.LMA (R)<br>Indicates IA-32e mode is active when set.  |   |
|                  |         | 11   | Execute Disable Bit Enable: IA32_EFER.NXE (R/W)   |   |
|                  |         | 63:12  | Reserved  |   |
| C000_0081H       |         | IA32_STAR  | System Call Target Address (R/W)  | If CPUID.80000001:EDX.[29] = 1  |
| C000_0082H       |         | IA32_LSTAR   | IA-32e Mode System Call Target Address (R/W)<br>Target RIP for the called procedure when SYSCALL is executed in 64-bit mode.  | If CPUID.80000001:EDX.[29] = 1  |
| C000_0083H       |         | IA32_CSTAR   | IA-32e Mode System Call Target Address (R/W)<br>Not used, as the SYSCALL instruction is not recognized in compatibility mode. | If CPUID.80000001:EDX.[29] = 1  |
| C000_0084H       |         | IA32_FMASK   | System Call Flag Mask (R/W)   | If CPUID.80000001:EDX.[29] = 1  |
| C000_0100H       |         | IA32_FS_BASE   | Map of BASE Address of FS (R/W)   | If CPUID.80000001:EDX.[29] = 1  |
| C000_0101H       |         | IA32_GS_BASE   | Map of BASE Address of GS (R/W)   | If CPUID.80000001:EDX.[29] = 1  |
| C000_0102H       |         | IA32_KERNEL_GS_BASE                                      | Swap Target of BASE Address of GS (R/W)   | If CPUID.80000001:EDX.[29] = 1  |
| C000_0103H       |         | IA32_TSC_AUX   | Auxiliary TSC (RW)  | If CPUID.80000001H: EDX[27] = 1 or CPUID.(EAX=7,ECX=0):ECX [bit 22] = 1 |
|                  |         | 31:0   | AUX: Auxiliary signature of TSC.  |   |
|                  |         | 63:32  | Reserved  |   |

**NOTES:**

1. In processors based on Intel NetBurst® microarchitecture, MSR addresses 180H-197H are supported, software must treat them as model-specific. Starting with Intel Core Duo processors, MSR addresses 180H-185H, 188H-197H are reserved.
2. The \*\_ADDR MSRs may or may not be present; this depends on flag settings in IA32\_MCI\_STATUS. See Section 15.3.2.3 and Section 15.3.2.4 for more information.
3. MAXPHYADDR is reported by CPUID.80000008H:EAX[7:0].

## 2.2 MSRS IN THE INTEL® CORE™ 2 PROCESSOR FAMILY

Table 2-3 lists model-specific registers (MSRs) for Intel Core 2 processor family and for Intel Xeon processors based on Intel Core microarchitecture, architectural MSR addresses are also included in Table 2-3. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_0FH, see Table 2-1.

MSRs listed in Table 2-2 and Table 2-3 are also supported by processors based on the Enhanced Intel Core microarchitecture. Processors based on the Enhanced Intel Core microarchitecture have the CPUID signature DisplayFamily\_DisplayModel of 06\_17H.

The column “Shared/Unique” applies to multi-core processors based on Intel Core microarchitecture. “Unique” means each processor core has a separate MSR, or a bit field in an MSR governs only a core independently. “Shared” means the MSR or the bit field in an MSR address governs the operation of both processor cores.

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description  |
|------------------|-----|----------------------------|----------------|--|
| Hex              | Dec |                            |                |  |
| 0H               | 0   | IA32_P5_MC_ADDR            | Unique         | See Section 2.22, “MSRs in Pentium Processors.”  |
| 1H               | 1   | IA32_P5_MC_TYPE            | Unique         | See Section 2.22, “MSRs in Pentium Processors.”  |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Unique         | See Section 8.10.5, “Monitor/Mwait Address Range Determination.” and Table 2-2.  |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Unique         | See Section 17.17, “Time-Stamp Counter,” and see Table 2-2.  |
| 17H              | 23  | IA32_PLATFORM_ID           | Shared         | Platform ID (R)<br>See Table 2-2.  |
| 17H              | 23  | MSR_PLATFORM_ID            | Shared         | Model Specific Platform ID (R)   |
|                  |     | 7:0                        |                | Reserved   |
|                  |     | 12:8                       |                | Maximum Qualified Ratio (R)<br>The maximum allowed bus ratio.  |
|                  |     | 49:13                      |                | Reserved   |
|                  |     | 52:50                      |                | See Table 2-2.   |
|                  |     | 63:53                      |                | Reserved   |
| 1BH              | 27  | IA32_APIC_BASE             | Unique         | See Section 10.4.4, “Local APIC Status and Location” and Table 2-2.  |
| 2AH              | 42  | MSR_EBL_CR_POWERON         | Shared         | Processor Hard Power-On Configuration (R/W)<br>Enables and disables processor features; (R) indicates current processor configuration. |
|                  |     | 0                          |                | Reserved   |
|                  |     | 1                          |                | Data Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processors implement R/W.                               |
|                  |     | 2                          |                | Response Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                           |

Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
|                  |     | 3                          |                   | MCERR# Drive Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processors implement R/W.   |
|                  |     | 4                          |                   | Address Parity Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processors implement R/W.   |
|                  |     | 5                          |                   | Reserved  |
|                  |     | 6                          |                   | Reserved  |
|                  |     | 7                          |                   | BINIT# Driver Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processors implement R/W.  |
|                  |     | 8                          |                   | Output Tri-state Enabled (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |     | 9                          |                   | Execute BIST (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |     | 10                         |                   | MCERR# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |     | 11                         |                   | Intel TXT Capable Chipset. (R/O)<br>1 = Present; 0 = Not Present  |
|                  |     | 12                         |                   | BINIT# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |     | 13                         |                   | Reserved  |
|                  |     | 14                         |                   | 1 MByte Power on Reset Vector (R/O)<br>1 = 1 MByte; 0 = 4 GBytes  |
|                  |     | 15                         |                   | Reserved  |
|                  |     | 17:16                      |                   | APIC Cluster ID (R/O)   |
|                  |     | 18                         |                   | N/2 Non-Integer Bus Ratio (R/O)<br>0 = Integer ratio; 1 = Non-integer ratio   |
|                  |     | 19                         |                   | Reserved  |
|                  |     | 21:20                      |                   | Symmetric Arbitration ID (R/O)  |
|                  |     | 26:22                      |                   | Integer Bus Frequency Ratio (R/O)   |
| 3AH              | 58  | MSR_FEATURE_CONTROL        | Unique            | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2.  |
|                  |     | 3                          | Unique            | SMRR Enable (R/WL)<br>When this bit is set and the lock bit is set, this makes the SMRR_PHYS_BASE and SMRR_PHYS_MASK registers read visible and writeable while in SMM. |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
| 40H              | 64  | MSR_LASTBRANCH_0_FROM_IP   | Unique            | Last Branch Record 0 From IP (R/W)<br>One of four pairs of last branch record registers on the last branch record stack. The From_IP part of the stack contains pointers to the source instruction. See also:<br><ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.5.</li> </ul> |
| 41H              | 65  | MSR_LASTBRANCH_1_FROM_IP   | Unique            | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 42H              | 66  | MSR_LASTBRANCH_2_FROM_IP   | Unique            | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 43H              | 67  | MSR_LASTBRANCH_3_FROM_IP   | Unique            | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 60H              | 96  | MSR_LASTBRANCH_0_TO_IP     | Unique            | Last Branch Record 0 To IP (R/W)<br>One of four pairs of last branch record registers on the last branch record stack. This To_IP part of the stack contains pointers to the destination instruction.   |
| 61H              | 97  | MSR_LASTBRANCH_1_TO_IP     | Unique            | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 62H              | 98  | MSR_LASTBRANCH_2_TO_IP     | Unique            | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 63H              | 99  | MSR_LASTBRANCH_3_TO_IP     | Unique            | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Unique            | BIOS Update Trigger Register (w)<br>See Table 2-2.  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | Unique            | BIOS Update Signature ID (RO)<br>See Table 2-2.   |
| A0H              | 160 | MSR_SMRR_PHYSBASE          | Unique            | System Management Mode Base Address register (wO in SMM)<br>Model-specific implementation of SMRR-like interface, read visible and write only in SMM.   |
|                  |     | 11:0                       |                   | Reserved  |
|                  |     | 31:12                      |                   | PhysBase: SMRR physical Base Address.   |
|                  |     | 63:32                      |                   | Reserved  |
| A1H              | 161 | MSR_SMRR_PHYSMASK          | Unique            | System Management Mode Physical Address Mask register (wO in SMM)<br>Model-specific implementation of SMRR-like interface, read visible and write only in SMM.  |
|                  |     | 10:0                       |                   | Reserved  |
|                  |     | 11                         |                   | Valid: Physical address base and range mask are valid.  |
|                  |     | 31:12                      |                   | PhysMask: SMRR physical address range mask.   |

Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|-----|----------------------------|-------------------|--|
| Hex              | Dec |                            |                   |  |
|                  |     | 63:32                      |                   | Reserved   |
| C1H              | 193 | IA32_PMC0                  | Unique            | Performance Counter Register<br>See Table 2-2.   |
| C2H              | 194 | IA32_PMC1                  | Unique            | Performance Counter Register<br>See Table 2-2.   |
| CDH              | 205 | MSR_FSB_FREQ               | Shared            | Scaleable Bus Speed(R0)<br>This field indicates the intended scaleable bus clock speed for processors based on Intel Core microarchitecture.   |
|                  |     | 2:0                        |                   | <ul style="list-style-type: none"> <li>▪ 101B: 100 MHz (FSB 400)</li> <li>▪ 001B: 133 MHz (FSB 533)</li> <li>▪ 011B: 167 MHz (FSB 667)</li> <li>▪ 010B: 200 MHz (FSB 800)</li> <li>▪ 000B: 267 MHz (FSB 1067)</li> <li>▪ 100B: 333 MHz (FSB 1333)</li> </ul>                                     |
|                  |     |                            |                   | 133.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 001B.<br>166.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 011B.   |
|                  |     |                            |                   | 266.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 000B.<br>333.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 100B.   |
|                  |     | 63:3                       |                   | Reserved   |
| CDH              | 205 | MSR_FSB_FREQ               | Shared            | Scaleable Bus Speed(R0)<br>This field indicates the intended scaleable bus clock speed for processors based on Enhanced Intel Core microarchitecture.  |
|                  |     | 2:0                        |                   | <ul style="list-style-type: none"> <li>▪ 101B: 100 MHz (FSB 400)</li> <li>▪ 001B: 133 MHz (FSB 533)</li> <li>▪ 011B: 167 MHz (FSB 667)</li> <li>▪ 010B: 200 MHz (FSB 800)</li> <li>▪ 000B: 267 MHz (FSB 1067)</li> <li>▪ 100B: 333 MHz (FSB 1333)</li> <li>▪ 110B: 400 MHz (FSB 1600)</li> </ul> |
|                  |     |                            |                   | 133.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 001B.<br>166.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 011B.   |
|                  |     |                            |                   | 266.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 110B.<br>333.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 111B.   |
|                  |     | 63:3                       |                   | Reserved   |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|-----|----------------------------|-------------------|--|
| Hex              | Dec |                            |                   |  |
| E7H              | 231 | IA32_MPERF                 | Unique            | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.   |
| E8H              | 232 | IA32_APERF                 | Unique            | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP               | Unique            | See Table 2-2.   |
|                  |     | 11                         | Unique            | SMRR Capability Using MSR 0A0H and 0A1H (R)  |
| 174H             | 372 | IA32_SYSENTER_CS           | Unique            | See Table 2-2.   |
| 175H             | 373 | IA32_SYSENTER_ESP          | Unique            | See Table 2-2.   |
| 176H             | 374 | IA32_SYSENTER_EIP          | Unique            | See Table 2-2.   |
| 179H             | 377 | IA32_MCG_CAP               | Unique            | See Table 2-2.   |
| 17AH             | 378 | IA32_MCG_STATUS            | Unique            | Global Machine Check Status  |
|                  |     | 0                          |                   | RIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If cleared, the program cannot be reliably restarted.                     |
|                  |     | 1                          |                   | EIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |
|                  |     | 2                          |                   | MCIP<br>When set, bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception. |
|                  |     | 63:3                       |                   | Reserved   |
| 186H             | 390 | IA32_PERFEVTSELO           | Unique            | See Table 2-2.   |
| 187H             | 391 | IA32_PERFEVTSEL1           | Unique            | See Table 2-2.   |
| 198H             | 408 | IA32_PERF_STATUS           | Shared            | See Table 2-2.   |
| 198H             | 408 | MSR_PERF_STATUS            | Shared            | Current performance status. See Section 14.1.1, "Software Interface For Initiating Performance State Transitions".   |
|                  |     | 15:0                       |                   | Current Performance State Value  |
|                  |     | 30:16                      |                   | Reserved   |
|                  |     | 31                         |                   | XE Operation (R/O).<br>If set, XE operation is enabled. Default is cleared.  |
|                  |     | 39:32                      |                   | Reserved   |



Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
|                  |     | 44:40                      |                   | Maximum Bus Ratio (R/O)<br>Indicates maximum bus ratio configured for the processor.  |
|                  |     | 45                         |                   | Reserved  |
|                  |     | 46                         |                   | Non-Integer Bus Ratio (R/O)<br>Indicates non-integer bus ratio is enabled. Applies processors based on Enhanced Intel Core microarchitecture.   |
|                  |     | 63:47                      |                   | Reserved  |
| 199H             | 409 | IA32_PERF_CTL              | Unique            | See Table 2-2.  |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Unique            | Clock Modulation (R/W)<br>See Table 2-2.<br>IA32_CLOCK_MODULATION MSR was originally named IA32_THERM_CONTROL MSR.  |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Unique            | Thermal Interrupt Control (R/W)<br>See Table 2-2.   |
| 19CH             | 412 | IA32_THERM_STATUS          | Unique            | Thermal Monitor Status (R/W)<br>See Table 2-2.  |
| 19DH             | 413 | MSR_THERM2_CTL             | Unique            | Thermal Monitor 2 Control   |
|                  |     | 15:0                       |                   | Reserved  |
|                  |     | 16                         |                   | TM_SELECT (R/W)<br>Mode of automatic thermal monitor:<br>0 = Thermal Monitor 1 (thermally-initiated on-die modulation of the stop-clock duty cycle).<br>1 = Thermal Monitor 2 (thermally-initiated frequency transitions).<br>If bit 3 of the IA32_MISC_ENABLE register is cleared, TM_SELECT has no effect. Neither TM1 nor TM2 are enabled. |
|                  |     | 63:16                      |                   | Reserved  |
| 1A0H             | 416 | IA32_MISC_ENABLE           |                   | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |     | 0                          |                   | Fast-Strings Enable<br>See Table 2-2.   |
|                  |     | 2:1                        |                   | Reserved  |
|                  |     | 3                          | Unique            | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2.  |
|                  |     | 6:4                        |                   | Reserved  |
|                  |     | 7                          | Shared            | Performance Monitoring Available (R)<br>See Table 2-2.  |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description  |
|------------------|-----|----------------------------|----------------|--|
| Hex              | Dec |                            |                |  |
|                  |     | 8                          |                | Reserved   |
|                  |     | 9                          |                | Hardware Prefetcher Disable (R/W)<br>When set, disables the hardware prefetcher operation on streams of data. When clear (default), enables the prefetch queue.<br>Disabling of the hardware prefetcher may impact processor performance.  |
|                  |     | 10                         | Shared         | FERR# Multiplexing Enable (R/W)<br>1 = FERR# asserted by the processor to indicate a pending break event within the processor.<br>0 = Indicates compatible FERR# signaling behavior.<br>This bit must be set to 1 to support XAPIC interrupt model usage.  |
|                  |     | 11                         | Shared         | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.  |
|                  |     | 12                         | Shared         | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.  |
|                  |     | 13                         | Shared         | TM2 Enable (R/W)<br>When this bit is set (1) and the thermal sensor indicates that the die temperature is at the pre-determined threshold, the Thermal Monitor 2 mechanism is engaged. TM2 will reduce the bus to core ratio and voltage according to the value last written to MSR_THERM2_CTL bits 15:0.  |
|                  |     |                            |                | When this bit is clear (0, default), the processor does not change the VID signals or the bus to core ratio when the processor enters a thermally managed state.<br>The BIOS must enable this feature if the TM2 feature flag (CPUID.1:ECX[8]) is set; if the TM2 feature flag is not set, this feature is not supported and BIOS must not alter the contents of the TM2 bit location.<br>The processor is operating out of specification if both this bit and the TM1 bit are set to 0. |
|                  |     | 15:14                      |                | Reserved   |
|                  |     | 16                         | Shared         | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.   |
|                  |     | 18                         | Shared         | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.   |

Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
|                  |     | 19                         | Shared            | <p>Adjacent Cache Line Prefetch Disable (R/W)</p> <p>When set to 1, the processor fetches the cache line that contains data currently required by the processor. When set to 0, the processor fetches cache lines that comprise a cache line pair (128 bytes).</p> <p>Single processor platforms should not set this bit. Server platforms should set or clear this bit based on platform performance observed in validation and testing.</p> <p>BIOS may contain a setup option that controls the setting of this bit.</p> |
|                  |     | 20                         | Shared            | <p>Enhanced Intel SpeedStep Technology Select Lock (R/WO)</p> <p>When set, this bit causes the following bits to become read-only:</p> <ul style="list-style-type: none"> <li>▪ Enhanced Intel SpeedStep Technology Select Lock (this bit).</li> <li>▪ Enhanced Intel SpeedStep Technology Enable bit.</li> </ul> <p>The bit must be set before an Enhanced Intel SpeedStep Technology transition is requested. This bit is cleared on reset.</p>   |
|                  |     | 21                         |                   | Reserved  |
|                  |     | 22                         | Shared            | <p>Limit CPUID Maxval (R/W)</p> <p>See Table 2-2.</p>   |
|                  |     | 23                         | Shared            | <p>xTPR Message Disable (R/W)</p> <p>See Table 2-2.</p>   |
|                  |     | 33:24                      |                   | Reserved  |
|                  |     | 34                         | Unique            | <p>XD Bit Disable (R/W)</p> <p>See Table 2-2.</p>   |
|                  |     | 36:35                      |                   | Reserved  |
|                  |     | 37                         | Unique            | <p>DCU Prefetcher Disable (R/W)</p> <p>When set to 1, the DCU L1 data cache prefetcher is disabled. The default value after reset is 0. BIOS may write '1' to disable this feature.</p> <p>The DCU prefetcher is an L1 data cache prefetcher. When the DCU prefetcher detects multiple loads from the same line done within a time limit, the DCU prefetcher assumes the next line will be required. The next line is prefetched in to the L1 data cache from memory or L2.</p>   |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
|                  |     | 38                         | Shared            | <p>IDA Disable (R/W)</p> <p>When set to 1 on processors that support IDA, the Intel Dynamic Acceleration feature (IDA) is disabled and the IDA_Enable feature flag will be cleared (CPUID.06H: EAX[1]=0).</p> <p>When set to a 0 on processors that support IDA, CPUID.06H: EAX[1] reports the processor's support of IDA is enabled.</p> <p>Note: The power-on default value is used by BIOS to detect hardware support of IDA. If the power-on default value is 1, IDA is available in the processor. If the power-on default value is 0, IDA is not available.</p> |
|                  |     | 39                         | Unique            | <p>IP Prefetcher Disable (R/W)</p> <p>When set to 1, the IP prefetcher is disabled. The default value after reset is 0. BIOS may write '1' to disable this feature.</p> <p>The IP prefetcher is an L1 data cache prefetcher. The IP prefetcher looks for sequential load history to determine whether to prefetch the next expected data into the L1 cache from memory or L2.</p>   |
|                  |     | 63:40                      |                   | Reserved  |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS         | Unique            | <p>Last Branch Record Stack TOS (R/W)</p> <p>Contains an index (bits 0-3) that points to the MSR containing the most recent branch record.</p> <p>See MSR_LASTBRANCH_0_FROM_IP (at 40H).</p>  |
| 1D9H             | 473 | IA32_DEBUGCTL              | Unique            | <p>Debug Control (R/W)</p> <p>See Table 2-2.</p>  |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Unique            | <p>Last Exception Record From Linear IP (R)</p> <p>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.</p>   |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Unique            | <p>Last Exception Record To Linear IP (R)</p> <p>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.</p>   |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Unique            | See Table 2-2.  |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Unique            | See Table 2-2.  |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Unique            | See Table 2-2.  |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Unique            | See Table 2-2.  |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Unique            | See Table 2-2.  |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Unique            | See Table 2-2.  |

Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Unique            | See Table 2-2.  |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Unique            | See Table 2-2.  |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Unique            | See Table 2-2.  |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Unique            | See Table 2-2.  |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Unique            | See Table 2-2.  |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Unique            | See Table 2-2.  |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Unique            | See Table 2-2.  |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6        | Unique            | See Table 2-2.  |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7        | Unique            | See Table 2-2.  |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7        | Unique            | See Table 2-2.  |
| 250H             | 592 | IA32_MTRR_FIX64K_00000     | Unique            | See Table 2-2.  |
| 258H             | 600 | IA32_MTRR_FIX16K_80000     | Unique            | See Table 2-2.  |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000     | Unique            | See Table 2-2.  |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000      | Unique            | See Table 2-2.  |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000      | Unique            | See Table 2-2.  |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000      | Unique            | See Table 2-2.  |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000      | Unique            | See Table 2-2.  |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000      | Unique            | See Table 2-2.  |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000      | Unique            | See Table 2-2.  |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000      | Unique            | See Table 2-2.  |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000      | Unique            | See Table 2-2.  |
| 277H             | 631 | IA32_PAT                   | Unique            | See Table 2-2.  |
| 2FFH             | 767 | IA32_MTRR_DEF_TYPE         | Unique            | Default Memory Types (R/W)<br>See Table 2-2.                          |
| 309H             | 777 | IA32_FIXED_CTR0            | Unique            | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2. |
| 309H             | 777 | MSR_PERF_FIXED_CTR0        | Unique            | Fixed-Function Performance Counter Register 0 (R/W)                   |
| 30AH             | 778 | IA32_FIXED_CTR1            | Unique            | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2. |
| 30AH             | 778 | MSR_PERF_FIXED_CTR1        | Unique            | Fixed-Function Performance Counter Register 1 (R/W)                   |
| 30BH             | 779 | IA32_FIXED_CTR2            | Unique            | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2. |
| 30BH             | 779 | MSR_PERF_FIXED_CTR2        | Unique            | Fixed-Function Performance Counter Register 2 (R/W)                   |
| 345H             | 837 | IA32_PERF_CAPABILITIES     | Unique            | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."               |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|------|----------------------------|-------------------|--|
| Hex              | Dec  |                            |                   |  |
| 345H             | 837  | MSR_PERF_CAPABILITIES      | Unique            | RO. This applies to processors that do not support architectural perfmon version 2.  |
|                  |      | 5:0                        |                   | LBR Format. See Table 2-2.   |
|                  |      | 6                          |                   | PEBS Record Format   |
|                  |      | 7                          |                   | PEBSSaveArchRegs. See Table 2-2.   |
|                  |      | 63:8                       |                   | Reserved   |
| 38DH             | 909  | IA32_FIXED_CTR_CTRL        | Unique            | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.  |
| 38DH             | 909  | MSR_PERF_FIXED_CTR_CTRL    | Unique            | Fixed-Function-Counter Control Register (R/W)  |
| 38EH             | 910  | IA32_PERF_GLOBAL_STATUS    | Unique            | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 38EH             | 910  | MSR_PERF_GLOBAL_STATUS     | Unique            | See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 38FH             | 911  | IA32_PERF_GLOBAL_CTRL      | Unique            | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 38FH             | 911  | MSR_PERF_GLOBAL_CTRL       | Unique            | See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 390H             | 912  | IA32_PERF_GLOBAL_OVF_CTRL  | Unique            | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 390H             | 912  | MSR_PERF_GLOBAL_OVF_CTRL   | Unique            | See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Unique            | See Table 2-2. See Section 18.6.2.4, "Processor Event Based Sampling (PEBS)."  |
|                  |      | 0                          |                   | Enable PEBS on IA32_PMC0. (R/W)  |
| 400H             | 1024 | IA32_MCO_CTL               | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 401H             | 1025 | IA32_MCO_STATUS            | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."  |
| 402H             | 1026 | IA32_MCO_ADDR              | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDRv flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 404H             | 1028 | IA32_MC1_CTL               | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 405H             | 1029 | IA32_MC1_STATUS            | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."  |
| 406H             | 1030 | IA32_MC1_ADDR              | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC1_ADDR register is either not implemented or contains no address if the ADDRv flag in the IA32_MC1_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 408H             | 1032 | IA32_MC2_CTL               | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |

Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|------|----------------------------|-------------------|---|
| Hex              | Dec  |                            |                   |   |
| 409H             | 1033 | IA32_MC2_STATUS            | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40CH             | 1036 | IA32_MC4_CTL               | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 40DH             | 1037 | IA32_MC4_STATUS            | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40EH             | 1038 | IA32_MC4_ADDR              | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 410H             | 1040 | IA32_MC3_CTL               |                   | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC3_STATUS            |                   | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC3_ADDR              | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 413H             | 1043 | IA32_MC3_MISC              | Unique            | Machine Check Error Reporting Register: Contains additional information describing the machine-check error if the MISC_V flag in the IA32_MCi_STATUS register is set.   |
| 414H             | 1044 | IA32_MC5_CTL               | Unique            | Machine Check Error Reporting Register: Controls signaling of #MC for errors produced by a particular hardware unit (or group of hardware units).   |
| 415H             | 1045 | IA32_MC5_STATUS            | Unique            | Machine Check Error Reporting Register: Contains information related to a machine-check error if its VAL (valid) flag is set. Software is responsible for clearing IA32_MCi_STATUS MSRs by explicitly writing 0s to them; writing 1s to them causes a general-protection exception.                           |
| 416H             | 1046 | IA32_MC5_ADDR              | Unique            | Machine Check Error Reporting Register: Contains the address of the code or data memory location that produced the machine-check error if the ADDR_V flag in the IA32_MCi_STATUS register is set.   |
| 417H             | 1047 | IA32_MC5_MISC              | Unique            | Machine Check Error Reporting Register: Contains additional information describing the machine-check error if the MISC_V flag in the IA32_MCi_STATUS register is set.   |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|------|----------------------------|-------------------|--|
| Hex              | Dec  |                            |                   |  |
| 419H             | 1045 | IA32_MC6_STATUS            | Unique            | Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 23. |
| 480H             | 1152 | IA32_VMX_BASIC             | Unique            | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."                         |
| 481H             | 1153 | IA32_VMX_PINBASED_CTLs     | Unique            | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls."     |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTLs    | Unique            | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."         |
| 483H             | 1155 | IA32_VMX_EXIT_CTLs         | Unique            | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Table 2-2.<br>See Appendix A.4, "VM-Exit Controls."                         |
| 484H             | 1156 | IA32_VMX_ENTRY_CTLs        | Unique            | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Table 2-2.<br>See Appendix A.5, "VM-Entry Controls."                       |
| 485H             | 1157 | IA32_VMX_MISC              | Unique            | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.6, "Miscellaneous Data."                    |
| 486H             | 1158 | IA32_VMX_CR0_FIXED0        | Unique            | Capability Reporting Register of CR0 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."                 |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1        | Unique            | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."                 |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0        | Unique            | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."                 |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1        | Unique            | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."                 |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM         | Unique            | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.<br>See Appendix A.9, "VMCS Enumeration."                   |



Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|------|----------------------------|-------------------|---|
| Hex              | Dec  |                            |                   |   |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL52   | Unique            | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."            |
| 600H             | 1536 | IA32_DS_AREA               | Unique            | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."   |
| 107CC<br>H       |      | MSR_EMON_L3_CTR_CTL0       | Unique            | GBUSQ Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2  |
| 107CD<br>H       |      | MSR_EMON_L3_CTR_CTL1       | Unique            | GBUSQ Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2  |
| 107CE<br>H       |      | MSR_EMON_L3_CTR_CTL2       | Unique            | GSPNPQ Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2 |
| 107CF<br>H       |      | MSR_EMON_L3_CTR_CTL3       | Unique            | GSPNPQ Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2 |
| 107D0<br>H       |      | MSR_EMON_L3_CTR_CTL4       | Unique            | FSB Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2    |
| 107D1<br>H       |      | MSR_EMON_L3_CTR_CTL5       | Unique            | FSB Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2    |
| 107D2<br>H       |      | MSR_EMON_L3_CTR_CTL6       | Unique            | FSB Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2    |
| 107D3<br>H       |      | MSR_EMON_L3_CTR_CTL7       | Unique            | FSB Event Control/Counter Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2    |
| 107D8<br>H       |      | MSR_EMON_L3_GL_CTL         | Unique            | L3/FSB Common Control Register (R/W)<br>Applies to Intel Xeon processor 7400 series (processor signature 06_1D) only. See Section 17.2.2        |
| C000_<br>0080H   |      | IA32_EFER                  | Unique            | Extended Feature Enables<br>See Table 2-2.  |
| C000_<br>0081H   |      | IA32_STAR                  | Unique            | System Call Target Address (R/W)<br>See Table 2-2.  |
| C000_<br>0082H   |      | IA32_LSTAR                 | Unique            | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.  |

**Table 2-3. MSRs in Processors Based on Intel® Core™ Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description   |
|------------------|-----|----------------------------|----------------|---|
| Hex              | Dec |                            |                |   |
| C000_0084H       |     | IA32_FMASK                 | Unique         | System Call Flag Mask (R/W)<br>See Table 2-2.             |
| C000_0100H       |     | IA32_FS_BASE               | Unique         | Map of BASE Address of FS (R/W)<br>See Table 2-2.         |
| C000_0101H       |     | IA32_GS_BASE               | Unique         | Map of BASE Address of GS (R/W)<br>See Table 2-2.         |
| C000_0102H       |     | IA32_KERNEL_GS_BASE        | Unique         | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2. |

### 2.3 MSRS IN THE 45 NM AND 32 NM INTEL® ATOM™ PROCESSOR FAMILY

Table 2-4 lists model-specific registers (MSRs) for 45 nm and 32 nm Intel Atom processors, architectural MSR addresses are also included in Table 2-4. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_1CH, 06\_26H, 06\_27H, 06\_35H and 06\_36H; see Table 2-1.

The column “Shared/Unique” applies to logical processors sharing the same core in processors based on the Intel Atom microarchitecture. “Unique” means each logical processor has a separate MSR, or a bit field in an MSR governs only a logical processor. “Shared” means the MSR or the bit field in an MSR address governs the operation of both logical processors in the same core.

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description   |
|------------------|-----|----------------------------|----------------|---|
| Hex              | Dec |                            |                |   |
| 0H               | 0   | IA32_P5_MC_ADDR            | Shared         | See Section 2.22, “MSRs in Pentium Processors.”                                 |
| 1H               | 1   | IA32_P5_MC_TYPE            | Shared         | See Section 2.22, “MSRs in Pentium Processors.”                                 |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Unique         | See Section 8.10.5, “Monitor/Mwait Address Range Determination.” and Table 2-2. |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Unique         | See Section 17.17, “Time-Stamp Counter,” and see Table 2-2.                     |
| 17H              | 23  | IA32_PLATFORM_ID           | Shared         | Platform ID (R)<br>See Table 2-2.   |
| 17H              | 23  | MSR_PLATFORM_ID            | Shared         | Model Specific Platform ID (R)  |
|                  |     | 7:0                        |                | Reserved  |
|                  |     | 12:8                       |                | Maximum Qualified Ratio (R)<br>The maximum allowed bus ratio.                   |
|                  |     | 63:13                      |                | Reserved  |
| 1BH              | 27  | IA32_APIC_BASE             | Unique         | See Section 10.4.4, “Local APIC Status and Location” and Table 2-2.             |

Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)

| Register Address |   | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|---|----------------------------|-------------------|--|
| Hex              | Dec   |                            |                   |  |
| 2AH              | 42  | MSR_EBL_CR_POWERON         | Shared            | Processor Hard Power-On Configuration (R/W)<br>Enables and disables processor features; (R) indicates current processor configuration. |
|                  |   | 0                          |                   | Reserved   |
|                  |   | 1                          |                   | Data Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Always 0.   |
|                  |   | 2                          |                   | Response Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Always 0.   |
|                  |   | 3                          |                   | AERR# Drive Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Always 0.   |
|                  |   | 4                          |                   | BERR# Enable for initiator bus requests (R/W)<br>1 = Enabled; 0 = Disabled<br>Always 0.  |
|                  |   | 5                          |                   | Reserved   |
|                  |   | 6                          |                   | Reserved   |
|                  |   | 7                          |                   | BINIT# Driver Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Always 0.   |
|                  |   | 8                          |                   | Reserved   |
|                  |   | 9                          |                   | Execute BIST (R/O)<br>1 = Enabled; 0 = Disabled  |
|                  |   | 10                         |                   | AERR# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled<br>Always 0.  |
|                  |   | 11                         |                   | Reserved   |
|                  |   | 12                         |                   | BINIT# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled<br>Always 0.   |
|                  |   | 13                         |                   | Reserved   |
|                  |   | 14                         |                   | 1 MByte Power on Reset Vector (R/O)<br>1 = 1 MByte; 0 = 4 GBytes   |
|                  |   | 15                         |                   | Reserved   |
|                  |   | 17:16                      |                   | APIC Cluster ID (R/O)<br>Always 00B.   |
| 19:18            | Reserved                                      |                            |                   |  |
| 21:20            | Symmetric Arbitration ID (R/O)<br>Always 00B. |                            |                   |  |

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description   |
|------------------|-----|----------------------------|----------------|---|
| Hex              | Dec |                            |                |   |
|                  |     | 26:22                      |                | Integer Bus Frequency Ratio (R/O)   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Unique         | Control Features in Intel 64Processor (R/W)<br>See Table 2-2.   |
| 40H              | 64  | MSR_LASTBRANCH_0_FROM_IP   | Unique         | Last Branch Record 0 From IP (R/W)<br>One of eight pairs of last branch record registers on the last branch record stack. The From_IP part of the stack contains pointers to the source instruction . See also:<br><ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.5.</li> </ul> |
| 41H              | 65  | MSR_LASTBRANCH_1_FROM_IP   | Unique         | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 42H              | 66  | MSR_LASTBRANCH_2_FROM_IP   | Unique         | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 43H              | 67  | MSR_LASTBRANCH_3_FROM_IP   | Unique         | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 44H              | 68  | MSR_LASTBRANCH_4_FROM_IP   | Unique         | Last Branch Record 4 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 45H              | 69  | MSR_LASTBRANCH_5_FROM_IP   | Unique         | Last Branch Record 5 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 46H              | 70  | MSR_LASTBRANCH_6_FROM_IP   | Unique         | Last Branch Record 6 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 47H              | 71  | MSR_LASTBRANCH_7_FROM_IP   | Unique         | Last Branch Record 7 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 60H              | 96  | MSR_LASTBRANCH_0_TO_IP     | Unique         | Last Branch Record 0 To IP (R/W)<br>One of eight pairs of last branch record registers on the last branch record stack. The To_IP part of the stack contains pointers to the destination instruction.   |
| 61H              | 97  | MSR_LASTBRANCH_1_TO_IP     | Unique         | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 62H              | 98  | MSR_LASTBRANCH_2_TO_IP     | Unique         | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 63H              | 99  | MSR_LASTBRANCH_3_TO_IP     | Unique         | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 64H              | 100 | MSR_LASTBRANCH_4_TO_IP     | Unique         | Last Branch Record 4 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 65H              | 101 | MSR_LASTBRANCH_5_TO_IP     | Unique         | Last Branch Record 5 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 66H              | 102 | MSR_LASTBRANCH_6_TO_IP     | Unique         | Last Branch Record 6 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 67H              | 103 | MSR_LASTBRANCH_7_TO_IP     | Unique         | Last Branch Record 7 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |

Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Shared            | BIOS Update Trigger Register (W)<br>See Table 2-2.  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | Unique            | BIOS Update Signature ID (RO)<br>See Table 2-2.   |
| C1H              | 193 | IA32_PMC0                  | Unique            | Performance counter register<br>See Table 2-2.  |
| C2H              | 194 | IA32_PMC1                  | Unique            | Performance Counter Register<br>See Table 2-2.  |
| CDH              | 205 | MSR_FSB_FREQ               | Shared            | Scaleable Bus Speed(RO)<br>This field indicates the intended scaleable bus clock speed for processors based on Intel Atom microarchitecture.  |
|                  |     | 2:0                        |                   | <ul style="list-style-type: none"> <li>▪ 111B: 083 MHz (FSB 333)</li> <li>▪ 101B: 100 MHz (FSB 400)</li> <li>▪ 001B: 133 MHz (FSB 533)</li> <li>▪ 011B: 167 MHz (FSB 667)</li> </ul> 133.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 001B.<br>166.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 011B. |
|                  |     | 63:3                       |                   | Reserved  |
| E7H              | 231 | IA32_MPERF                 | Unique            | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| E8H              | 232 | IA32_APERF                 | Unique            | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.   |
| FEH              | 254 | IA32_MTRRCAP               | Shared            | Memory Type Range Register (R)<br>See Table 2-2.  |
| 11EH             | 281 | MSR_BBL_CR_CTL3            | Shared            | Control Register 3<br>Used to configure the L2 Cache.   |
|                  |     | 0                          |                   | L2 Hardware Enabled (RO)<br>1 = Indicates the L2 is hardware-enabled.<br>0 = Indicates the L2 is hardware-disabled.   |
|                  |     | 7:1                        |                   | Reserved  |
|                  |     | 8                          |                   | L2 Enabled (R/W)<br>1 = L2 cache has been initialized.<br>0 = Disabled (default).<br>Until this bit is set, the processor will not respond to the WBINVD instruction or the assertion of the FLUSH# input.  |
|                  |     | 22:9                       |                   | Reserved  |
|                  |     | 23                         |                   | L2 Not Present (RO)<br>0 = L2 Present<br>1 = L2 Not Present   |
|                  |     | 63:24                      |                   | Reserved  |

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|-----|----------------------------|-------------------|--|
| Hex              | Dec |                            |                   |  |
| 174H             | 372 | IA32_SYSENTER_CS           | Unique            | See Table 2-2.   |
| 175H             | 373 | IA32_SYSENTER_ESP          | Unique            | See Table 2-2.   |
| 176H             | 374 | IA32_SYSENTER_EIP          | Unique            | See Table 2-2.   |
| 179H             | 377 | IA32_MCG_CAP               | Unique            | See Table 2-2.   |
| 17AH             | 378 | IA32_MCG_STATUS            | Unique            | Global Machine Check Status  |
|                  |     | 0                          |                   | RIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If cleared, the program cannot be reliably restarted.                     |
|                  |     | 1                          |                   | EIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |
|                  |     | 2                          |                   | MCIP<br>When set, bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception. |
|                  |     | 63:3                       |                   | Reserved   |
| 186H             | 390 | IA32_PERFEVTSELO           | Unique            | See Table 2-2.   |
| 187H             | 391 | IA32_PERFEVTSEL1           | Unique            | See Table 2-2.   |
| 198H             | 408 | IA32_PERF_STATUS           | Shared            | See Table 2-2.   |
| 198H             | 408 | MSR_PERF_STATUS            | Shared            | Performance Status   |
|                  |     | 15:0                       |                   | Current Performance State Value  |
|                  |     | 39:16                      |                   | Reserved   |
|                  |     | 44:40                      |                   | Maximum Bus Ratio (R/O)<br>Indicates maximum bus ratio configured for the processor.   |
|                  |     | 63:45                      |                   | Reserved   |
| 199H             | 409 | IA32_PERF_CTL              | Unique            | See Table 2-2.   |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Unique            | Clock Modulation (R/W)<br>See Table 2-2.<br>IA32_CLOCK_MODULATION MSR was originally named IA32_THERM_CONTROL MSR.   |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Unique            | Thermal Interrupt Control (R/W)<br>See Table 2-2.  |
| 19CH             | 412 | IA32_THERM_STATUS          | Unique            | Thermal Monitor Status (R/W)<br>See Table 2-2.   |
| 19DH             | 413 | MSR_THERM2_CTL             | Shared            | Thermal Monitor 2 Control  |

Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------------|-------------------|---|
| Hex              | Dec |                            |                   |   |
|                  |     | 15:0                       |                   | Reserved  |
|                  |     | 16                         |                   | TM_SELECT (R/W)<br>Mode of automatic thermal monitor:<br>0 = Thermal Monitor 1 (thermally-initiated on-die modulation of the stop-clock duty cycle).<br>1 = Thermal Monitor 2 (thermally-initiated frequency transitions).<br>If bit 3 of the IA32_MISC_ENABLE register is cleared, TM_SELECT has no effect. Neither TM1 nor TM2 are enabled. |
|                  |     | 63:17                      |                   | Reserved  |
| 1A0H             | 416 | IA32_MISC_ENABLE           | Unique            | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |     | 0                          |                   | Fast-Strings Enable<br>See Table 2-2.   |
|                  |     | 2:1                        |                   | Reserved  |
|                  |     | 3                          | Unique            | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2. Default value is 0.  |
|                  |     | 6:4                        |                   | Reserved  |
|                  |     | 7                          | Shared            | Performance Monitoring Available (R)<br>See Table 2-2.  |
|                  |     | 8                          |                   | Reserved  |
|                  |     | 9                          |                   | Reserved  |
|                  |     | 10                         | Shared            | FERR# Multiplexing Enable (R/W)<br>1 = FERR# asserted by the processor to indicate a pending break event within the processor.<br>0 = Indicates compatible FERR# signaling behavior.<br>This bit must be set to 1 to support XAPIC interrupt model usage.   |
|                  |     | 11                         | Shared            | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.   |
|                  |     | 12                         | Shared            | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.   |
|                  |     | 13                         | Shared            | TM2 Enable (R/W)<br>When this bit is set (1) and the thermal sensor indicates that the die temperature is at the pre-determined threshold, the Thermal Monitor 2 mechanism is engaged. TM2 will reduce the bus to core ratio and voltage according to the value last written to MSR_THERM2_CTL bits 15:0.                                     |

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Shared/ Unique | Bit Description   |
|------------------|-----|----------------------------|----------------|---|
| Hex              | Dec |                            |                |   |
|                  |     |                            |                | <p>When this bit is cleared (0, default), the processor does not change the VID signals or the bus to core ratio when the processor enters a thermally managed state.</p> <p>The BIOS must enable this feature if the TM2 feature flag (CPUID.1:ECX[8]) is set; if the TM2 feature flag is not set, this feature is not supported and BIOS must not alter the contents of the TM2 bit location.</p> <p>The processor is operating out of specification if both this bit and the TM1 bit are set to 0.</p> |
|                  |     | 15:14                      |                | Reserved  |
|                  |     | 16                         | Shared         | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.  |
|                  |     | 18                         | Shared         | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.  |
|                  |     | 19                         |                | Reserved  |
|                  |     | 20                         | Shared         | <p>Enhanced Intel SpeedStep Technology Select Lock (R/WO)</p> <p>When set, this bit causes the following bits to become read-only:</p> <ul style="list-style-type: none"> <li>▪ Enhanced Intel SpeedStep Technology Select Lock (this bit).</li> <li>▪ Enhanced Intel SpeedStep Technology Enable bit.</li> </ul> <p>The bit must be set before an Enhanced Intel SpeedStep Technology transition is requested. This bit is cleared on reset.</p>   |
|                  |     | 21                         |                | Reserved  |
|                  |     | 22                         | Unique         | Limit CPUID Maxval (R/W)<br>See Table 2-2.  |
|                  |     | 23                         | Shared         | xTPR Message Disable (R/W)<br>See Table 2-2.  |
|                  |     | 33:24                      |                | Reserved  |
|                  |     | 34                         | Unique         | XD Bit Disable (R/W)<br>See Table 2-2.  |
|                  |     | 63:35                      |                | Reserved  |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS         | Unique         | <p>Last Branch Record Stack TOS (R/W)</p> <p>Contains an index (bits 0-2) that points to the MSR containing the most recent branch record.<br/>See MSR_LASTBRANCH_0_FROM_IP (at 40H).</p>   |
| 1D9H             | 473 | IA32_DEBUGCTL              | Unique         | <p>Debug Control (R/W)</p> <p>See Table 2-2.</p>  |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Unique         | <p>Last Exception Record From Linear IP (R)</p> <p>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.</p>   |



Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Shared/<br>Unique | Bit Description  |
|------------------|-----|----------------------------|-------------------|--|
| Hex              | Dec |                            |                   |  |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Unique            | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled. |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Shared            | See Table 2-2.   |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Shared            | See Table 2-2.   |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Shared            | See Table 2-2.   |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Shared            | See Table 2-2.   |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Shared            | See Table 2-2.   |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Shared            | See Table 2-2.   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Shared            | See Table 2-2.   |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Shared            | See Table 2-2.   |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Shared            | See Table 2-2.   |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Shared            | See Table 2-2.   |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Shared            | See Table 2-2.   |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Shared            | See Table 2-2.   |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Shared            | See Table 2-2.   |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6        | Shared            | See Table 2-2.   |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7        | Shared            | See Table 2-2.   |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7        | Shared            | See Table 2-2.   |
| 250H             | 592 | IA32_MTRR_FIX64K_00000     | Shared            | See Table 2-2.   |
| 258H             | 600 | IA32_MTRR_FIX16K_80000     | Shared            | See Table 2-2.   |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000     | Shared            | See Table 2-2.   |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000      | Shared            | See Table 2-2.   |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000      | Shared            | See Table 2-2.   |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000      | Shared            | See Table 2-2.   |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000      | Shared            | See Table 2-2.   |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000      | Shared            | See Table 2-2.   |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000      | Shared            | See Table 2-2.   |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000      | Shared            | See Table 2-2.   |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000      | Shared            | See Table 2-2.   |
| 277H             | 631 | IA32_PAT                   | Unique            | See Table 2-2.   |
| 309H             | 777 | IA32_FIXED_CTR0            | Unique            | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2.  |
| 30AH             | 778 | IA32_FIXED_CTR1            | Unique            | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2.  |

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)**

| Register Address |      | Register Name / Bit Fields | Shared/ Unique | Bit Description   |
|------------------|------|----------------------------|----------------|---|
| Hex              | Dec  |                            |                |   |
| 30BH             | 779  | IA32_FIXED_CTR2            | Unique         | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2.   |
| 345H             | 837  | IA32_PERF_CAPABILITIES     | Shared         | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."   |
| 38DH             | 909  | IA32_FIXED_CTR_CTRL        | Unique         | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.   |
| 38EH             | 910  | IA32_PERF_GLOBAL_STATUS    | Unique         | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 38FH             | 911  | IA32_PERF_GLOBAL_CTRL      | Unique         | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 390H             | 912  | IA32_PERF_GLOBAL_OVF_CTRL  | Unique         | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Unique         | See Table 2-2. See Section 18.6.2.4, "Processor Event Based Sampling (PEBS)."   |
|                  |      | 0                          |                |   |
| 400H             | 1024 | IA32_MCO_CTL               | Shared         | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 401H             | 1025 | IA32_MCO_STATUS            | Shared         | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 402H             | 1026 | IA32_MCO_ADDR              | Shared         | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 404H             | 1028 | IA32_MC1_CTL               | Shared         | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS            | Shared         | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 408H             | 1032 | IA32_MC2_CTL               | Shared         | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 409H             | 1033 | IA32_MC2_STATUS            | Shared         | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Shared         | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40CH             | 1036 | IA32_MC3_CTL               | Shared         | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 40DH             | 1037 | IA32_MC3_STATUS            | Shared         | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 40EH             | 1038 | IA32_MC3_ADDR              | Shared         | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |

Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)

| Register Address |      | Register Name / Bit Fields | Shared/<br>Unique | Bit Description   |
|------------------|------|----------------------------|-------------------|---|
| Hex              | Dec  |                            |                   |   |
| 410H             | 1040 | IA32_MC4_CTL               | Shared            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC4_STATUS            | Shared            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC4_ADDR              | Shared            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 480H             | 1152 | IA32_VMX_BASIC             | Unique            | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."  |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL      | Unique            | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls."  |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL     | Unique            | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."  |
| 483H             | 1155 | IA32_VMX_EXIT_CTL          | Unique            | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Table 2-2.<br>See Appendix A.4, "VM-Exit Controls."  |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL         | Unique            | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Table 2-2.<br>See Appendix A.5, "VM-Entry Controls."  |
| 485H             | 1157 | IA32_VMX_MISC              | Unique            | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.6, "Miscellaneous Data."   |
| 486H             | 1158 | IA32_VMX_CR0_FIXED0        | Unique            | Capability Reporting Register of CR0 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."  |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1        | Unique            | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."  |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0        | Unique            | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."  |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1        | Unique            | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."  |

**Table 2-4. MSRs in 45 nm and 32 nm Intel® Atom™ Processor Family (Contd.)**

| Register Address |      | Register Name / Bit Fields | Shared/ Unique | Bit Description  |
|------------------|------|----------------------------|----------------|--|
| Hex              | Dec  |                            |                |  |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM         | Unique         | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.<br>See Appendix A.9, "VMCS Enumeration."             |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTLSD   | Unique         | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls." |
| 600H             | 1536 | IA32_DS_AREA               | Unique         | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."  |
| C000_0080H       |      | IA32_EFER                  | Unique         | Extended Feature Enables<br>See Table 2-2.   |
| C000_0081H       |      | IA32_STAR                  | Unique         | System Call Target Address (R/W)<br>See Table 2-2.   |
| C000_0082H       |      | IA32_LSTAR                 | Unique         | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.   |
| C000_0084H       |      | IA32_FMASK                 | Unique         | System Call Flag Mask (R/W)<br>See Table 2-2.  |
| C000_0100H       |      | IA32_FS_BASE               | Unique         | Map of BASE Address of FS (R/W)<br>See Table 2-2.  |
| C000_0101H       |      | IA32_GS_BASE               | Unique         | Map of BASE Address of GS (R/W)<br>See Table 2-2.  |
| C000_0102H       |      | IA32_KERNEL_GS_BASE        | Unique         | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.  |

Table 2-5 lists model-specific registers (MSRs) that are specific to Intel® Atom™ processor with the CPUID signature with DisplayFamily\_DisplayModel of 06\_27H.

**Table 2-5. MSRs Supported by Intel® Atom™ Processors with CPUID Signature 06\_27H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 3F8H             | 1016 | MSR_PKG_C2_RESIDENCY       | Package | Package C2 Residency<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 63:0                       | Package | Package C2 Residency Counter (R/O)<br>Time that this package is in processor-specific C2 states since last reset. Counts at 1 Mhz frequency.              |

**Table 2-5. MSRs Supported by Intel® Atom™ Processors (Contd.)with CPUID Signature 06\_27H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 3F9H             | 1017 | MSR_PKG_C4_RESIDENCY       | Package | Package C4 Residency<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 63:0                       | Package | Package C4 Residency Counter. (R/O)<br>Time that this package is in processor-specific C4 states since last reset. Counts at 1 Mhz frequency.             |
| 3FAH             | 1018 | MSR_PKG_C6_RESIDENCY       | Package | Package C6 Residency<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 63:0                       | Package | Package C6 Residency Counter. (R/O)<br>Time that this package is in processor-specific C6 states since last reset. Counts at 1 Mhz frequency.             |

## 2.4 MSRS IN INTEL PROCESSORS BASED ON SILVERMONT MICROARCHITECTURE

Table 2-6 lists model-specific registers (MSRs) common to Intel processors based on the Silvermont microarchitecture. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_37H, 06\_4AH, 06\_4DH, 06\_5AH, and 06\_5DH; see Table 2-1. The MSRs listed in Table 2-6 are also common to processors based on the Airmont microarchitecture and newer microarchitectures for next generation Intel Atom processors.

Table 2-7 lists MSRs common to processors based on the Silvermont and Airmont microarchitectures, but not newer microarchitectures.

Table 2-8, Table 2-9, and Table 2-10 lists MSRs that are model-specific across processors based on the Silvermont microarchitecture.

In the Silvermont microarchitecture, the scope column indicates the following: "Core" means each processor core has a separate MSR, or a bit field not shared with another processor core. "Module" means the MSR or the bit field is shared by a pair of processor cores in the physical package. "Package" means all processor cores in the physical package share the same MSR or bit interface.

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
| 0H               | 0   | IA32_P5_MC_ADDR            | Module | See Section 2.22, "MSRs in Pentium Processors."                                 |
| 1H               | 1   | IA32_P5_MC_TYPE            | Module | See Section 2.22, "MSRs in Pentium Processors."                                 |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Core   | See Section 8.10.5, "Monitor/Mwait Address Range Determination," and Table 2-2. |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Core   | See Section 17.17, "Time-Stamp Counter," and Table 2-2.                         |
| 1BH              | 27  | IA32_APIC_BASE             | Core   | See Section 10.4.4, "Local APIC Status and Location," and Table 2-2.            |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 2AH              | 42  | MSR_EBL_CR_POWERON         | Module | Processor Hard Power-On Configuration (R/W)<br>Writes ignored.   |
|                  |     | 63:0                       |        | Reserved   |
| 34H              | 52  | MSR_SMI_COUNT              | Core   | SMI Counter (R/O)  |
|                  |     | 31:0                       |        | SMI Count (R/O)<br>Running count of SMI events since last RESET.   |
|                  |     | 63:32                      |        | Reserved   |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Core   | BIOS Update Trigger Register (W)<br>See Table 2-2.   |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | Core   | BIOS Update Signature ID (RO)<br>See Table 2-2.  |
| C1H              | 193 | IA32_PMC0                  | Core   | Performance counter register<br>See Table 2-2.   |
| C2H              | 194 | IA32_PMC1                  | Core   | Performance Counter Register<br>See Table 2-2.   |
| E4H              | 228 | MSR_PMG_IO_CAPTURE_BASE    | Module | Power Management IO Redirection in C-state (R/W)<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 15:0                       |        | LVL_2 Base Address (R/W)<br>Specifies the base address visible to software for IO redirection. If IO MWAIT Redirection is enabled, reads to this address will be consumed by the power management logic and decoded to MWAIT instructions. When IO port address redirection is enabled, this is the IO port address reported to the OS/software. |
|                  |     | 18:16                      |        | C-state Range (R/W)<br>Specifies the encoding value of the maximum C-State code name to be included when IO read to MWAIT redirection is enabled by MSR_PKG_CST_CONFIG_CONTROL[bit10]:<br>100b - C4 is the max C-State to include<br>110b - C6 is the max C-State to include<br>111b - C7 is the max C-State to include                          |
|                  |     | 63:19                      |        | Reserved   |
| E7H              | 231 | IA32_MPERF                 | Core   | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.   |
| E8H              | 232 | IA32_APERF                 | Core   | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP               | Core   | Memory Type Range Register (R)<br>See Table 2-2.   |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
| 13CH             | 52  | MSR_FEATURE_CONFIG         | Core  | AES Configuration (RW-L)<br>Privileged post-BIOS agent must provide a #GP handler to handle unsuccessful read of this MSR.  |
|                  |     | 1:0                        |       | AES Configuration (RW-L)<br>Upon a successful read of this MSR, the configuration of AES instruction sets availability is as follows:<br>11b: AES instructions are not available until next RESET.<br>Otherwise, AES instructions are available.<br>Note: AES instruction set is not available if read is unsuccessful. If the configuration is not 01b, AES instructions can be mis-configured if a privileged agent unintentionally writes 11b. |
|                  |     | 63:2                       |       | Reserved  |
| 174H             | 372 | IA32_SYSENTER_CS           | Core  | See Table 2-2.  |
| 175H             | 373 | IA32_SYSENTER_ESP          | Core  | See Table 2-2.  |
| 176H             | 374 | IA32_SYSENTER_EIP          | Core  | See Table 2-2.  |
| 179H             | 377 | IA32_MCG_CAP               | Core  | See Table 2-2.  |
| 17AH             | 378 | IA32_MCG_STATUS            | Core  | Global Machine Check Status   |
|                  |     | 0                          |       | RIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If cleared, the program cannot be reliably restarted.  |
|                  |     | 1                          |       | EIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.   |
|                  |     | 2                          |       | MCIP<br>When set, bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception.  |
|                  |     | 63:3                       |       | Reserved  |
| 186H             | 390 | IA32_PERFVTSELO            | Core  | See Table 2-2.  |
|                  |     | 7:0                        |       | Event Select  |
|                  |     | 15:8                       |       | UMask   |
|                  |     | 16                         |       | USR   |
|                  |     | 17                         |       | OS  |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 18                         |         | Edge   |
|                  |     | 19                         |         | PC   |
|                  |     | 20                         |         | INT  |
|                  |     | 21                         |         | Reserved   |
|                  |     | 22                         |         | EN   |
|                  |     | 23                         |         | INV  |
|                  |     | 31:24                      |         | CMASK  |
|                  |     | 63:32                      |         | Reserved   |
| 187H             | 391 | IA32_PERFEVTSEL1           | Core    | See Table 2-2.   |
| 198H             | 408 | IA32_PERF_STATUS           | Module  | See Table 2-2.   |
| 199H             | 409 | IA32_PERF_CTL              | Core    | See Table 2-2.   |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Core    | Clock Modulation (R/W)<br>See Table 2-2.<br>IA32_CLOCK_MODULATION MSR was originally named IA32_THERM_CONTROL MSR.   |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Core    | Thermal Interrupt Control (R/W)<br>See Table 2-2.  |
| 19CH             | 412 | IA32_THERM_STATUS          | Core    | Thermal Monitor Status (R/W)<br>See Table 2-2.   |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Package | Temperature Target   |
|                  |     | 15:0                       |         | Reserved   |
|                  |     | 23:16                      |         | Temperature Target (R)<br>The default thermal throttling or PROCHOT# activation temperature in degrees C. The effective temperature for thermal throttling or PROCHOT# activation is "Temperature Target" + "Target Offset". |
|                  |     | 29:24                      |         | Target Offset (R/W)<br>Specifies an offset in degrees C to adjust the throttling and PROCHOT# activation temperature from the default target specified in TEMPERATURE_TARGET (bits 23:16).                                   |
|                  |     | 63:30                      |         | Reserved   |
| 1A6H             | 422 | MSR_OFFCORE_RSP_0          | Module  | Offcore Response Event Select Register (R/W)   |
| 1A7H             | 423 | MSR_OFFCORE_RSP_1          | Module  | Offcore Response Event Select Register (R/W)   |
| 1BOH             | 432 | IA32_ENERGY_PERF_BIAS      | Core    | See Table 2-2.   |
| 1D9H             | 473 | IA32_DEBUGCTL              | Core    | Debug Control (R/W)<br>See Table 2-2.  |



**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Core  | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.                       |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Core  | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled. |
| 1F2H             | 498 | IA32_SMRR_PHYSBASE         | Core  | See Table 2-2.   |
| 1F3H             | 499 | IA32_SMRR_PHYSMASK         | Core  | See Table 2-2.   |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Core  | See Table 2-2.   |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Core  | See Table 2-2.   |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Core  | See Table 2-2.   |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Core  | See Table 2-2.   |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Core  | See Table 2-2.   |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Core  | See Table 2-2.   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Core  | See Table 2-2.   |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Core  | See Table 2-2.   |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Core  | See Table 2-2.   |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Core  | See Table 2-2.   |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Core  | See Table 2-2.   |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Core  | See Table 2-2.   |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Core  | See Table 2-2.   |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6        | Core  | See Table 2-2.   |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7        | Core  | See Table 2-2.   |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7        | Core  | See Table 2-2.   |
| 250H             | 592 | IA32_MTRR_FIX64K_00000     | Core  | See Table 2-2.   |
| 258H             | 600 | IA32_MTRR_FIX16K_80000     | Core  | See Table 2-2.   |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000     | Core  | See Table 2-2.   |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000      | Core  | See Table 2-2.   |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000      | Core  | See Table 2-2.   |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000      | Core  | See Table 2-2.   |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000      | Core  | See Table 2-2.   |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000      | Core  | See Table 2-2.   |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000      | Core  | See Table 2-2.   |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000      | Core  | See Table 2-2.   |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000      | Core  | See Table 2-2.   |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 277H             | 631  | IA32_PAT                   | Core   | See Table 2-2.  |
| 2FFH             | 767  | IA32_MTRR_DEF_TYPE         | Core   | Default Memory Types (R/W)<br>See Table 2-2.  |
| 309H             | 777  | IA32_FIXED_CTR0            | Core   | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2.   |
| 30AH             | 778  | IA32_FIXED_CTR1            | Core   | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2.   |
| 30BH             | 779  | IA32_FIXED_CTR2            | Core   | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2.   |
| 345H             | 837  | IA32_PERF_CAPABILITIES     | Core   | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."   |
| 38DH             | 909  | IA32_FIXED_CTR_CTRL        | Core   | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.   |
| 38FH             | 911  | IA32_PERF_GLOBAL_CTRL      | Core   | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."   |
| 3FDH             | 1021 | MSR_CORE_C6_RESIDENCY      | Core   | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.   |
|                  |      | 63:0                       |        | CORE C6 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C6 states. Counts at the TSC Frequency.   |
| 400H             | 1024 | IA32_MCO_CTL               | Module | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 401H             | 1025 | IA32_MCO_STATUS            | Module | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 402H             | 1026 | IA32_MCO_ADDR              | Module | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 404H             | 1028 | IA32_MC1_CTL               | Module | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS            | Module | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 408H             | 1032 | IA32_MC2_CTL               | Module | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 409H             | 1033 | IA32_MC2_STATUS            | Module | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Module | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 40CH             | 1036 | IA32_MC3_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 40DH             | 1037 | IA32_MC3_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40EH             | 1038 | IA32_MC3_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 410H             | 1040 | IA32_MC4_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC4_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC4_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 415H             | 1045 | IA32_MC5_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 416H             | 1046 | IA32_MC5_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 480H             | 1152 | IA32_VMX_BASIC             | Core    | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."  |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL      | Core    | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls."  |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL     | Core    | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."  |
| 483H             | 1155 | IA32_VMX_EXIT_CTL          | Core    | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Table 2-2.<br>See Appendix A.4, "VM-Exit Controls."  |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |      | Register Name / Bit Fields   | Scope | Bit Description  |
|------------------|------|------------------------------|-------|--|
| Hex              | Dec  |                              |       |  |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL5          | Core  | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Table 2-2.<br>See Appendix A.5, "VM-Entry Controls."                 |
| 485H             | 1157 | IA32_VMX_MISC                | Core  | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.6, "Miscellaneous Data."              |
| 486H             | 1158 | IA32_VMX_CR0_FIXED0          | Core  | Capability Reporting Register of CR0 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."           |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1          | Core  | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."           |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0          | Core  | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."           |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1          | Core  | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."           |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM           | Core  | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.<br>See Appendix A.9, "VMCS Enumeration."             |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL52     | Core  | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls." |
| 48CH             | 1164 | IA32_VMX_EPT_VPID_ENUM       | Core  | Capability Reporting Register of EPT and VPID (R/O)<br>See Table 2-2   |
| 48DH             | 1165 | IA32_VMX_TRUE_PINBASED_CTL5  | Core  | Capability Reporting Register of Pin-Based VM-Execution Flex Controls (R/O)<br>See Table 2-2   |
| 48EH             | 1166 | IA32_VMX_TRUE_PROCBASED_CTL5 | Core  | Capability Reporting Register of Primary Processor-based VM-Execution Flex Controls (R/O)<br>See Table 2-2                           |
| 48FH             | 1167 | IA32_VMX_TRUE_EXIT_CTL5      | Core  | Capability Reporting Register of VM-Exit Flex Controls (R/O)<br>See Table 2-2  |

**Table 2-6. MSRs Common to the Silvermont Microarchitecture and Newer Microarchitectures for Intel Atom Processors**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|------|----------------------------|-------|--|
| Hex              | Dec  |                            |       |  |
| 490H             | 1168 | IA32_VMX_TRUE_ENTRY_CTL5   | Core  | Capability Reporting Register of VM-Entry Flex Controls (R/O)<br>See Table 2-2   |
| 491H             | 1169 | IA32_VMX_FMFUNC            | Core  | Capability Reporting Register of VM-Function Controls (R/O)<br>See Table 2-2   |
| 4C1H             | 1217 | IA32_A_PMC0                | Core  | See Table 2-2.   |
| 4C2H             | 1218 | IA32_A_PMC1                | Core  | See Table 2-2.   |
| 600H             | 1536 | IA32_DS_AREA               | Core  | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."  |
| 660H             | 1632 | MSR_CORE_C1_RESIDENCY      | Core  | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.          |
|                  |      | 63:0                       |       | CORE C1 Residency Counter. (R/O)<br>Value since last reset that this core is in processor-specific C1 states. Counts at the TSC frequency. |
| 6E0H             | 1760 | IA32_TSC_DEADLINE          | Core  | TSC Target of Local APIC's TSC Deadline Mode (R/W)<br>See Table 2-2.   |
| C000_0080H       |      | IA32_EFER                  | Core  | Extended Feature Enables<br>See Table 2-2.   |
| C000_0081H       |      | IA32_STAR                  | Core  | System Call Target Address (R/W)<br>See Table 2-2.   |
| C000_0082H       |      | IA32_LSTAR                 | Core  | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.   |
| C000_0084H       |      | IA32_FMASK                 | Core  | System Call Flag Mask (R/W)<br>See Table 2-2.  |
| C000_0100H       |      | IA32_FS_BASE               | Core  | Map of BASE Address of FS (R/W)<br>See Table 2-2.  |
| C000_0101H       |      | IA32_GS_BASE               | Core  | Map of BASE Address of GS (R/W)<br>See Table 2-2.  |
| C000_0102H       |      | IA32_KERNEL_GS_BASE        | Core  | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.  |
| C000_0103H       |      | IA32_TSC_AUX               | Core  | AUXILIARY TSC Signature (R/W)<br>See Table 2-2   |

Table 2-7 lists model-specific registers (MSRs) that are common to Intel® Atom™ processors based on the Silvermont and Airmont microarchitectures but not newer microarchitectures.

**Table 2-7. MSRs Common to the Silvermont and Airmont Microarchitectures**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
| 17H              | 23  | MSR_PLATFORM_ID            | Module | Model Specific Platform ID (R)  |
|                  |     | 7:0                        |        | Reserved  |
|                  |     | 13:8                       |        | Maximum Qualified Ratio (R)<br>The maximum allowed bus ratio.   |
|                  |     | 49:13                      |        | Reserved  |
|                  |     | 52:50                      |        | See Table 2-2.  |
|                  |     | 63:33                      |        | Reserved  |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Core   | Control Features in Intel 64Processor (R/W)<br>See Table 2-2.   |
|                  |     | 0                          |        | Lock (R/WL)   |
|                  |     | 1                          |        | Reserved  |
|                  |     | 2                          |        | Enable VMX outside SMX operation (R/WL)   |
| 40H              | 64  | MSR_LASTBRANCH_0_FROM_IP   | Core   | Last Branch Record 0 From IP (R/W)<br>One of eight pairs of last branch record registers on the last branch record stack. The From_IP part of the stack contains pointers to the source instruction. See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.5 and record format in Section 17.4.8.1.</li> </ul> |
| 41H              | 65  | MSR_LASTBRANCH_1_FROM_IP   | Core   | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 42H              | 66  | MSR_LASTBRANCH_2_FROM_IP   | Core   | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 43H              | 67  | MSR_LASTBRANCH_3_FROM_IP   | Core   | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 44H              | 68  | MSR_LASTBRANCH_4_FROM_IP   | Core   | Last Branch Record 4 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 45H              | 69  | MSR_LASTBRANCH_5_FROM_IP   | Core   | Last Branch Record 5 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 46H              | 70  | MSR_LASTBRANCH_6_FROM_IP   | Core   | Last Branch Record 6 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 47H              | 71  | MSR_LASTBRANCH_7_FROM_IP   | Core   | Last Branch Record 7 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 60H              | 96  | MSR_LASTBRANCH_0_TO_IP     | Core   | Last Branch Record 0 To IP (R/W)<br>One of eight pairs of last branch record registers on the last branch record stack. The To_IP part of the stack contains pointers to the destination instruction.   |
| 61H              | 97  | MSR_LASTBRANCH_1_TO_IP     | Core   | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |

Table 2-7. MSRs Common to the Silvermont and Airmont Microarchitectures

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 62H              | 98  | MSR_LASTBRANCH_2_TO_IP     | Core   | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 63H              | 99  | MSR_LASTBRANCH_3_TO_IP     | Core   | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 64H              | 100 | MSR_LASTBRANCH_4_TO_IP     | Core   | Last Branch Record 4 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 65H              | 101 | MSR_LASTBRANCH_5_TO_IP     | Core   | Last Branch Record 5 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 66H              | 102 | MSR_LASTBRANCH_6_TO_IP     | Core   | Last Branch Record 6 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 67H              | 103 | MSR_LASTBRANCH_7_TO_IP     | Core   | Last Branch Record 7 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Module | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0                        |        | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0 (no package C-state support)<br>001b: C1 (Behavior is the same as 000b)<br>100b: C4<br>110b: C6<br>111b: C7 (Silvermont only). |
|                  |     | 9:3                        |        | Reserved   |
|                  |     | 10                         |        | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.  |
|                  |     | 14:11                      |        | Reserved   |
|                  |     | 15                         |        | CFG Lock (R/WO)<br>When set, locks bits 15:0 of this register until next reset.  |
|                  |     | 63:16                      |        | Reserved   |
| 11EH             | 281 | MSR_BBL_CR_CTL3            | Module | Control Register 3<br>Used to configure the L2 Cache.  |
|                  |     | 0                          |        | L2 Hardware Enabled (RO)<br>1 = If the L2 is hardware-enabled.<br>0 = Indicates if the L2 is hardware-disabled.  |

**Table 2-7. MSRs Common to the Silvermont and Airmont Microarchitectures**

| Register Address |        | Register Name / Bit Fields                   | Scope  | Bit Description   |
|------------------|--------|--|--------|---|
| Hex              | Dec    |  |        |   |
|                  |        | 7:1  |        | Reserved  |
|                  |        | 8  |        | L2 Enabled (R/W)<br>1 = L2 cache has been initialized.<br>0 = Disabled (default).<br>Until this bit is set the processor will not respond to the WBINVD instruction or the assertion of the FLUSH# input. |
|                  |        | 22:9   |        | Reserved  |
|                  |        | 23   |        | L2 Not Present (RO)<br>0 = L2 Present.<br>1 = L2 Not Present.   |
|                  |        | 63:24  |        | Reserved  |
| 1A0H             | 416    | IA32_MISC_ENABLE                             |        | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |        | 0  | Core   | Fast-Strings Enable<br>See Table 2-2.   |
|                  |        | 2:1  |        | Reserved  |
|                  |        | 3  | Module | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2. Default value is 0.  |
|                  |        | 6:4  |        | Reserved  |
|                  |        | 7  | Core   | Performance Monitoring Available (R)<br>See Table 2-2.  |
|                  |        | 10:8   |        | Reserved  |
|                  |        | 11   | Core   | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.   |
|                  |        | 12   | Core   | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.   |
|                  |        | 15:13  |        | Reserved  |
|                  |        | 16   | Module | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.  |
|                  |        | 18   | Core   | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.  |
|                  |        | 21:19  |        | Reserved  |
|                  |        | 22   | Core   | Limit CPUID Maxval (R/W)<br>See Table 2-2.  |
| 23               | Module | xTPR Message Disable (R/W)<br>See Table 2-2. |        |   |
| 33:24            |        | Reserved                                     |        |   |



Table 2-7. MSRs Common to the Silvermont and Airmont Microarchitectures

| Register Address |      | Register Name / Bit Fields | Scope    | Bit Description  |
|------------------|------|----------------------------|----------|--|
| Hex              | Dec  |                            |          |  |
|                  |      | 34                         | Core     | XD Bit Disable (R/W)<br>See Table 2-2.   |
|                  |      | 37:35                      |          | Reserved   |
|                  |      | 38                         | Module   | Turbo Mode Disable (R/W)<br>When set to 1 on processors that support Intel Turbo Boost Technology, the turbo mode feature is disabled and the IDA_Enable feature flag will be cleared (CPUID.06H: EAX[1]=0).<br>When set to a 0 on processors that support IDA, CPUID.06H: EAX[1] reports the processor's support of turbo mode is enabled.<br>Note: The power-on default value is used by BIOS to detect hardware support of turbo mode. If the power-on default value is 1, turbo mode is available in the processor. If the power-on default value is 0, turbo mode is not available. |
|                  |      | 63:39                      |          | Reserved   |
| 1C8H             | 456  | MSR_LBR_SELECT             | Core     | Last Branch Record Filtering Select Register (R/W)<br>See Section 17.9.2, "Filtering of Last Branch Records."  |
|                  |      | 0                          |          | CPL_EQ_0   |
|                  |      | 1                          |          | CPL_NEQ_0  |
|                  |      | 2                          |          | JCC  |
|                  |      | 3                          |          | NEAR_REL_CALL  |
|                  |      | 4                          |          | NEAR_IND_CALL  |
|                  |      | 5                          |          | NEAR_RET   |
|                  |      | 6                          |          | NEAR_IND_JMP   |
|                  |      | 7                          |          | NEAR_REL_JMP   |
|                  |      | 8                          |          | FAR_BRANCH   |
|                  | 63:9 |                            | Reserved |  |
| 1C9H             | 457  | MSR_LASTBRANCH_TOS         | Core     | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-2) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_0_FROM_IP.   |
| 38EH             | 910  | IA32_PERF_GLOBAL_STATUS    | Core     | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 390H             | 912  | IA32_PERF_GLOBAL_OVF_CTRL  | Core     | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Core     | See Table 2-2. See Section 18.6.2.4, "Processor Event Based Sampling (PEBS)."  |
|                  |      | 0                          |          | Enable PEBS for precise event on IA32_PMC0 (R/W)   |

**Table 2-7. MSRs Common to the Silvermont and Airmont Microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 3FAH             | 1018 | MSR_PKG_C6_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                                      |
|                  |      | 63:0                       |         | Package C6 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C6 states. Counts at the TSC Frequency.                        |
| 664H             | 1636 | MSR_MC6_RESIDENCY_COUNTER  | Module  | Module C6 Residency Counter (R/O)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 63:0                       |         | Time that this module is in module-specific C6 states since last reset. Counts at 1 Mhz frequency.   |
| CEH              | 206  | MSR_PLATFORM_INFO          | Package | Platform Information: Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .             |
|                  |      | 7:0                        |         | Reserved   |
|                  |      | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the maximum frequency that does not require turbo. Frequency = ratio * Scalable Bus Frequency.                   |
|                  |      | 63:16                      |         | Reserved   |

### 2.4.1 MSRs with Model-Specific Behavior in the Silvermont Microarchitecture

Table 2-8 lists model-specific registers (MSRs) that are specific to Intel® Atom™ processor E3000 Series (CPUID signature with DisplayFamily\_DisplayModel of 06\_37H) and Intel Atom processors (CPUID signatures with DisplayFamily\_DisplayModel of 06\_4AH, 06\_5AH, 06\_5DH).

**Table 2-8. Specific MSRs Supported by Intel® Atom™ Processors with CPUID Signatures 06\_37H, 06\_4AH, 06\_5AH, 06\_5DH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers used in RAPL Interfaces (R/O)<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 3:0                        |         | Power Units<br>Power related information (in milliwatts) is based on the multiplier, $2^{PU}$ ; where PU is an unsigned integer represented by bits 3:0. Default value is 0101b, indicating power unit is in 32 milliwatts increment. |
|                  |      | 7:4                        |         | Reserved  |

**Table 2-8. Specific MSRs Supported by Intel® Atom™ Processors with CPUID Signatures 06\_37H, 06\_4AH, 06\_5AH, 06\_5DH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 12:8                       |         | Energy Status Units<br>Energy related information (in microJoules) is based on the multiplier, 2 <sup>^</sup> ESU; where ESU is an unsigned integer represented by bits 12:8. Default value is 00101b, indicating energy unit is in 32 microJoules increment. |
|                  |      | 15:13                      |         | Reserved  |
|                  |      | 19:16                      |         | Time Unit<br>The value is 0000b, indicating time unit is in one second.   |
|                  |      | 63:20                      |         | Reserved  |
| 610H             | 1552 | MSR_PKG_POWER_LIMIT        | Package | PKG RAPL Power Limit Control (R/W)  |
|                  |      | 14:0                       |         | Package Power Limit #1 (R/W)<br>See Section 14.9.3, "Package RAPL Domain." and MSR_RAPL_POWER_UNIT in Table 2-8.  |
|                  |      | 15                         |         | Enable Power Limit #1 (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
|                  |      | 16                         |         | Package Clamping Limitation #1 (R/W)<br>See Section 14.9.3, "Package RAPL Domain."  |
|                  |      | 23:17                      |         | Time Window for Power Limit #1 (R/W)<br>In unit of second. If 0 is specified in bits [23:17], defaults to 1 second window.  |
|                  |      | 63:24                      |         | Reserved  |
| 611H             | 1553 | MSR_PKG_ENERGY_STATUS      | Package | PKG Energy Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain." and MSR_RAPL_POWER_UNIT in Table 2-8.   |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS      | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains." and MSR_RAPL_POWER_UNIT in Table 2-8.  |
| CDH              | 205  | MSR_FSB_FREQ               | Module  | Scaleable Bus Speed(R/O)<br>This field indicates the intended scaleable bus clock speed for processors based on Silvermont microarchitecture.   |
|                  |      | 2:0                        |         | <ul style="list-style-type: none"> <li>▪ 100B: 080.0 MHz</li> <li>▪ 000B: 083.3 MHz</li> <li>▪ 001B: 100.0 MHz</li> <li>▪ 010B: 133.3 MHz</li> <li>▪ 011B: 116.7 MHz</li> </ul>   |
|                  |      | 63:3                       |         | Reserved  |

Table 2-9 lists model-specific registers (MSRs) that are specific to Intel® Atom™ processor E3000 Series (CPUID signature with DisplayFamily\_DisplayModel of 06\_37H).

**Table 2-9. Specific MSRs Supported by Intel® Atom™ Processor E3000 Series with CPUID Signature 06\_37H**

| Register Address |      | Register Name / Bit Fields     | Scope   | Bit Description  |
|------------------|------|--------------------------------|---------|--|
| Hex              | Dec  |                                |         |  |
| 668H             | 1640 | MSR_CC6_DEMOTION_POLICY_CONFIG | Package | Core C6 Demotion Policy Config MSR   |
|                  |      | 63:0                           |         | Controls per-core C6 demotion policy. Writing a value of 0 disables core level HW demotion policy.   |
| 669H             | 1641 | MSR_MC6_DEMOTION_POLICY_CONFIG | Package | Module C6 Demotion Policy Config MSR   |
|                  |      | 63:0                           |         | Controls module (i.e., two cores sharing the second-level cache) C6 demotion policy. Writing a value of 0 disables module level HW demotion policy.                    |
| 664H             | 1636 | MSR_MC6_RESIDENCY_COUNTER      | Module  | Module C6 Residency Counter (R/O)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 63:0                           |         | Time that this module is in module-specific C6 states since last reset. Counts at 1 Mhz frequency.   |

Table 2-10 lists model-specific registers (MSRs) that are specific to Intel® Atom™ processor C2000 Series (CPUID signature with DisplayFamily\_DisplayModel of 06\_4DH).

**Table 2-10. Specific MSRs Supported by Intel® Atom™ Processor C2000 Series with CPUID Signature 06\_4DH**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 1A4H             | 420 | MSR_MISC_FEATURE_CONTROL   |         | Miscellaneous Feature Control (R/W)  |
|                  |     | 0                          | Core    | L2 Hardware Prefetcher Disable (R/W)<br>If 1, disables the L2 hardware prefetcher, which fetches additional lines of code or data into the L2 cache. |
|                  |     | 1                          |         | Reserved   |
|                  |     | 2                          | Core    | DCU Hardware Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache prefetcher, which fetches the next cache line into L1 data cache.          |
|                  |     | 63:3                       |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode (RW)   |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |

**Table 2-10. Specific MSRs Supported by Intel® Atom™ Processor C2000 Series (Contd.)with CPUID Signature**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 39:32                      | Package | Maximum Ratio Limit for 5C<br>Maximum turbo ratio limit of 5 core active.  |
|                  |      | 47:40                      | Package | Maximum Ratio Limit for 6C<br>Maximum turbo ratio limit of 6 core active.  |
|                  |      | 55:48                      | Package | Maximum Ratio Limit for 7C<br>Maximum turbo ratio limit of 7 core active.  |
|                  |      | 63:56                      | Package | Maximum Ratio Limit for 8C<br>Maximum turbo ratio limit of 8 core active.  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers used in RAPL Interfaces (R/O)<br>See Section 14.9.1, "RAPL Interfaces."   |
|                  |      | 3:0                        |         | Power Units<br>Power related information (in milliwatts) is based on the multiplier, $2^{PU}$ ; where PU is an unsigned integer represented by bits 3:0. Default value is 0101b, indicating power unit is in 32 milliwatts increment.                  |
|                  |      | 7:4                        |         | Reserved   |
|                  |      | 12:8                       |         | Energy Status Units.<br>Energy related information (in microjoules) is based on the multiplier, $2^{ESU}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 00101b, indicating energy unit is in 32 microjoules increment. |
|                  |      | 15:13                      |         | Reserved   |
|                  |      | 19:16                      |         | Time Unit<br>The value is 0000b, indicating time unit is in one second.  |
|                  |      | 63:20                      |         | Reserved   |
| 610H             | 1552 | MSR_PKG_POWER_LIMIT        | Package | PKG RAPL Power Limit Control (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 66EH             | 1646 | MSR_PKG_POWER_INFO         | Package | PKG RAPL Parameter (R/O)   |
|                  |      | 14:0                       |         | Thermal Spec Power (R/O)<br>The unsigned integer value is the equivalent of the thermal specification power of the package domain. The unit of this field is specified by the "Power Units" field of MSR_RAPL_POWER_UNIT.                              |
|                  |      | 63:15                      |         | Reserved   |

## 2.4.2 MSRs In Intel Atom Processors Based on Airmont Microarchitecture

Intel Atom processor X7-Z8000 and X5-Z8000 series are based on the Airmont microarchitecture. These processors support MSRs listed in Table 2-6, Table 2-7, Table 2-8, and Table 2-11. These processors have a CPUID signature with DisplayFamily\_DisplayModel including 06\_4CH; see Table 2-1.

**Table 2-11. MSRs in Intel Atom Processors Based on the Airmont Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
| CDH              | 205 | MSR_FSB_FREQ               | Module | Scaleable Bus Speed(RO)<br>This field indicates the intended scaleable bus clock speed for processors based on Airmont microarchitecture.   |
|                  |     | 3:0                        |        | <ul style="list-style-type: none"> <li>▪ 0000B: 083.3 MHz</li> <li>▪ 0001B: 100.0 MHz</li> <li>▪ 0010B: 133.3 MHz</li> <li>▪ 0011B: 116.7 MHz</li> <li>▪ 0100B: 080.0 MHz</li> <li>▪ 0101B: 093.3 MHz</li> <li>▪ 0110B: 090.0 MHz</li> <li>▪ 0111B: 088.9 MHz</li> <li>▪ 1000B: 087.5 MHz</li> </ul>                                |
|                  |     | 63:5                       |        | Reserved  |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Module | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 2:0                        |        | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: No limit<br>001b: C1<br>010b: C2<br>110b: C6<br>111b: C7 |
|                  |     | 9:3                        |        | Reserved  |
|                  |     | 10                         |        | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.   |
|                  |     | 14:11                      |        | Reserved  |
|                  |     | 15                         |        | CFG Lock (R/WO)<br>When set, locks bits 15:0 of this register until next reset.   |
|                  |     | 63:16                      |        | Reserved  |
| E4H              | 228 | MSR_PMG_IO_CAPTURE_BASE    | Module | Power Management IO Redirection in C-state (R/W)<br>See <a href="http://biosbits.org">http://biosbits.org</a> .   |

Table 2-11. MSRs in Intel Atom Processors Based on the Airmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 15:0                       |         | LVL_2 Base Address (R/W)<br>Specifies the base address visible to software for IO redirection. If IO MWAIT Redirection is enabled, reads to this address will be consumed by the power management logic and decoded to MWAIT instructions. When IO port address redirection is enabled, this is the IO port address reported to the OS/software.   |
|                  |      | 18:16                      |         | C-state Range (R/W)<br>Specifies the encoding value of the maximum C-State code name to be included when IO read to MWAIT redirection is enabled by MSR_PKG_CST_CONFIG_CONTROL[bit10]:<br>000b - C3 is the max C-State to include.<br>001b - Deep Power Down Technology is the max C-State.<br>010b - C7 is the max C-State to include.  |
|                  |      | 63:19                      |         | Reserved   |
| 638H             | 1592 | MSR_PPO_POWER_LIMIT        | Package | PPO RAPL Power Limit Control (R/W)   |
|                  |      | 14:0                       |         | PPO Power Limit #1 (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains" and MSR_RAPL_POWER_UNIT in Table 2-8.   |
|                  |      | 15                         |         | Enable Power Limit #1 (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."   |
|                  |      | 16                         |         | Reserved   |
|                  |      | 23:17                      |         | Time Window for Power Limit #1 (R/W)<br>Specifies the time duration over which the average power must remain below PPO_POWER_LIMIT #1(14:0). Supported Encodings:<br>0x0: 1 second time duration.<br>0x1: 5 second time duration (Default).<br>0x2: 10 second time duration.<br>0x3: 15 second time duration.<br>0x4: 20 second time duration.<br>0x5: 25 second time duration.<br>0x6: 30 second time duration.<br>0x7: 35 second time duration.<br>0x8: 40 second time duration.<br>0x9: 45 second time duration.<br>0xA: 50 second time duration.<br>0xB-0x7F - reserved. |
|                  |      | 63:24                      |         | Reserved   |

## 2.5 MSRS IN INTEL ATOM PROCESSORS BASED ON GOLDMONT MICROARCHITECTURE

Intel Atom processors based on the Goldmont microarchitecture support MSRs listed in Table 2-6 and Table 2-12. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_5CH; see Table 2-1.

In the Goldmont microarchitecture, the scope column indicates the following: "Core" means each processor core has a separate MSR, or a bit field not shared with another processor core. "Module" means the MSR or the bit field is shared by a pair of processor cores in the physical package. "Package" means all processor cores in the physical package share the same MSR or bit interface.

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 17H              | 23  | MSR_PLATFORM_ID            | Module  | Model Specific Platform ID (R)   |
|                  |     | 49:0                       |         | Reserved   |
|                  |     | 52:50                      |         | See Table 2-2.   |
|                  |     | 63:33                      |         | Reserved   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Core    | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2.   |
|                  |     | 0                          |         | Lock (R/WL)  |
|                  |     | 1                          |         | Enable VMX inside SMX operation (R/WL)   |
|                  |     | 2                          |         | Enable VMX outside SMX operation (R/WL)  |
|                  |     | 14:8                       |         | SENTER local functions enables (R/WL)  |
|                  |     | 15                         |         | SENTER global functions enable (R/WL)  |
|                  |     | 18                         |         | SGX global functions enable (R/WL)   |
|                  |     | 63:19                      |         | Reserved   |
| 3BH              | 59  | IA32_TSC_ADJUST            | Core    | Per-Core TSC ADJUST (R/W)<br>See Table 2-2.  |
| C3H              | 195 | IA32_PMC2                  | Core    | Performance Counter Register<br>See Table 2-2.   |
| C4H              | 196 | IA32_PMC3                  | Core    | Performance Counter Register<br>See Table 2-2.   |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> . |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the maximum frequency that does not require turbo. Frequency = ratio * 100 MHz.                        |
|                  |     | 27:16                      |         | Reserved   |



Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled.   |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates TDP Limit for Turbo mode is not programmable.  |
|                  |     | 30                         | Package | Programmable TJ OFFSET (R/O)<br>When set to 1, indicates that MSR_TEMPERATURE_TARGET.[27:24] is valid and writable to specify a temperature offset.  |
|                  |     | 39:31                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz.   |
|                  |     | 63:48                      |         | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 3:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br><br>The following C-state code name encodings are supported:<br>0000b: No limit<br>0001b: C1<br>0010b: C3<br>0011b: C6<br>0100b: C7<br>0101b: C7S<br>0110b: C8<br>0111b: C9<br>1000b: C10 |
|                  |     | 9:3                        |         | Reserved   |
|                  |     | 10                         |         | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.  |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields                 | Scope   | Bit Description   |
|------------------|------|--|---------|---|
| Hex              | Dec  |  |         |   |
|                  |      | 14:11                                      |         | Reserved  |
|                  |      | 15   |         | CFG Lock (R/WO)<br>When set, locks bits 15:0 of this register until next reset.   |
|                  |      | 63:16                                      |         | Reserved  |
| 17DH             | 381  | MSR_SMM_MCA_CAP                            | Core    | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability enhancement. Accessible only while in SMM.   |
|                  |      | 57:0                                       |         | Reserved  |
|                  |      | 58   |         | SMM_Code_Access_Chk (SMM-RO)<br>If set to 1 indicates that the SMM code access restriction is supported and the MSR_SMM_FEATURE_CONTROL is supported. |
|                  |      | 59   |         | Long_Flow_Indication (SMM-RO)<br>If set to 1 indicates that the SMM long flow indicator is supported and the MSR_SMM_DELAYED is supported.            |
|                  |      | 63:60                                      |         | Reserved  |
| 188H             | 392  | IA32_PERFEVTSEL2                           | Core    | See Table 2-2.  |
| 189H             | 393  | IA32_PERFEVTSEL3                           | Core    | See Table 2-2.  |
| 1A0H             | 416  | IA32_MISC_ENABLE                           |         | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |      | 0  | Core    | Fast-Strings Enable<br>See Table 2-2.   |
|                  |      | 2:1  |         | Reserved  |
|                  |      | 3  | Package | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2. Default value is 1.  |
|                  |      | 6:4  |         | Reserved  |
|                  |      | 7  | Core    | Performance Monitoring Available (R)<br>See Table 2-2.  |
|                  |      | 10:8                                       |         | Reserved  |
|                  |      | 11   | Core    | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.   |
|                  |      | 12   | Core    | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.   |
|                  |      | 15:13                                      |         | Reserved  |
|                  |      | 16   | Package | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.  |
| 18               | Core | ENABLE MONITOR FSM (R/W)<br>See Table 2-2. |         |   |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 21:19                      |         | Reserved   |
|                  |     | 22                         | Core    | Limit CPUID Maxval (R/W)<br>See Table 2-2.   |
|                  |     | 23                         | Package | xTPR Message Disable (R/W)<br>See Table 2-2.   |
|                  |     | 33:24                      |         | Reserved   |
|                  |     | 34                         | Core    | XD Bit Disable (R/W)<br>See Table 2-2.   |
|                  |     | 37:35                      |         | Reserved   |
|                  |     | 38                         | Package | Turbo Mode Disable (R/W)<br>When set to 1 on processors that support Intel Turbo Boost Technology, the turbo mode feature is disabled and the IDA_Enable feature flag will be clear (CPUID.06H: EAX[1]=0).<br>When set to a 0 on processors that support IDA, CPUID.06H: EAX[1] reports the processor's support of turbo mode is enabled.<br>Note: The power-on default value is used by BIOS to detect hardware support of turbo mode. If the power-on default value is 1, turbo mode is available in the processor. If the power-on default value is 0, turbo mode is not available. |
|                  |     | 63:39                      |         | Reserved   |
| 1A4H             | 420 | MSR_MISC_FEATURE_CONTROL   |         | Miscellaneous Feature Control (R/W)  |
|                  |     | 0                          | Core    | L2 Hardware Prefetcher Disable (R/W)<br>If 1, disables the L2 hardware prefetcher, which fetches additional lines of code or data into the L2 cache.   |
|                  |     | 1                          |         | Reserved   |
|                  |     | 2                          | Core    | DCU Hardware Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache prefetcher, which fetches the next cache line into L1 data cache.  |
|                  |     | 63:3                       |         | Reserved   |
| 1AAH             | 426 | MSR_MISC_PWR_MGMT          | Package | Miscellaneous Power Management Control<br>Various model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 0                          |         | EIST Hardware Coordination Disable (R/W)<br>When 0, enables hardware coordination of Enhanced Intel Speedstep Technology request from processor cores. When 1, disables hardware coordination of Enhanced Intel Speedstep Technology requests.   |
|                  |     | 21:1                       |         | Reserved   |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 22                         |         | Thermal Interrupt Coordination Enable (R/W)<br>If set, then thermal interrupt on one core is routed to all cores.   |
|                  |     | 63:23                      |         | Reserved  |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode by Core Groups (RW)<br>Specifies Maximum Ratio Limit for each Core Group. Max ratio for groups with more cores must decrease monotonically.<br>For groups with less than 4 cores, the max ratio must be 32 or less. For groups with 4-5 cores, the max ratio must be 22 or less. For groups with more than 5 cores, the max ratio must be 16 or less. |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for Active Cores in Group 0<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 0 threshold.  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for Active Cores in Group 1<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 1 threshold, and greater than the Group 0 threshold.  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for Active Cores in Group 2<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 2 threshold, and greater than the Group 1 threshold.  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for Active Cores in Group 3<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 3 threshold, and greater than the Group 2 threshold.  |
|                  |     | 39:32                      | Package | Maximum Ratio Limit for Active Cores in Group 4<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 4 threshold, and greater than the Group 3 threshold.  |
|                  |     | 47:40                      | Package | Maximum Ratio Limit for Active Cores in Group 5<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 5 threshold, and greater than the Group 4 threshold.  |
|                  |     | 55:48                      | Package | Maximum Ratio Limit for Active Cores in Group 6<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 6 threshold, and greater than the Group 5 threshold.  |
|                  |     | 63:56                      | Package | Maximum Ratio Limit for Active Cores in Group 7<br>Maximum turbo ratio limit when the number of active cores is less than or equal to the Group 7 threshold, and greater than the Group 6 threshold.  |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 1AEH             | 430 | MSR_TURBO_GROUP_CORECNT    | Package | Group Size of Active Cores for Turbo Mode Operation (R/W)<br>Writes of 0 threshold is ignored.   |
|                  |     | 7:0                        | Package | Group 0 Core Count Threshold<br>Maximum number of active cores to operate under the Group 0 Max Turbo Ratio limit.   |
|                  |     | 15:8                       | Package | Group 1 Core Count Threshold<br>Maximum number of active cores to operate under the Group 1 Max Turbo Ratio limit. Must be greater than the Group 0 Core Count.  |
|                  |     | 23:16                      | Package | Group 2 Core Count Threshold<br>Maximum number of active cores to operate under the Group 2 Max Turbo Ratio limit. Must be greater than the Group 1 Core Count.  |
|                  |     | 31:24                      | Package | Group 3 Core Count Threshold<br>Maximum number of active cores to operate under the Group 3 Max Turbo Ratio limit. Must be greater than the Group 2 Core Count.  |
|                  |     | 39:32                      | Package | Group 4 Core Count Threshold<br>Maximum number of active cores to operate under the Group 4 Max Turbo Ratio limit. Must be greater than the Group 3 Core Count.  |
|                  |     | 47:40                      | Package | Group 5 Core Count Threshold<br>Maximum number of active cores to operate under the Group 5 Max Turbo Ratio limit. Must be greater than the Group 4 Core Count.  |
|                  |     | 55:48                      | Package | Group 6 Core Count Threshold<br>Maximum number of active cores to operate under the Group 6 Max Turbo Ratio limit. Must be greater than the Group 5 Core Count.  |
|                  |     | 63:56                      | Package | Group 7 Core Count Threshold<br>Maximum number of active cores to operate under the Group 7 Max Turbo Ratio limit. Must be greater than the Group 6 Core Count, and not less than the total number of processor cores in the package. E.g., specify 255. |
| 1C8H             | 456 | MSR_LBR_SELECT             | Core    | Last Branch Record Filtering Select Register (R/W)<br>See Section 17.9.2, "Filtering of Last Branch Records."  |
|                  |     | 0                          |         | CPL_EQ_0   |
|                  |     | 1                          |         | CPL_NEQ_0  |
|                  |     | 2                          |         | JCC  |
|                  |     | 3                          |         | NEAR_REL_CALL  |
|                  |     | 4                          |         | NEAR_IND_CALL  |
|                  |     | 5                          |         | NEAR_RET   |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 6                          |         | NEAR_IND_JMP   |
|                  |     | 7                          |         | NEAR_REL_JMP   |
|                  |     | 8                          |         | FAR_BRANCH   |
|                  |     | 9                          |         | EN_CALL_STACK  |
|                  |     | 63:10                      |         | Reserved   |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS         | Core    | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-4) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_O_FROM_IP.             |
| 1FCH             | 508 | MSR_POWER_CTL              | Core    | Power Control Register. See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 0                          |         | Reserved   |
|                  |     | 1                          | Package | C1E Enable (R/W)<br>When set to '1', will enable the CPU to switch to the Minimum Enhanced Intel SpeedStep Technology operating point when all execution cores enter MWAIT (C1). |
|                  |     | 63:2                       |         | Reserved   |
| 210H             | 528 | IA32_MTRR_PHYSBASE8        | Core    | See Table 2-2.   |
| 211H             | 529 | IA32_MTRR_PHYSMASK8        | Core    | See Table 2-2.   |
| 212H             | 530 | IA32_MTRR_PHYSBASE9        | Core    | See Table 2-2.   |
| 213H             | 531 | IA32_MTRR_PHYSMASK9        | Core    | See Table 2-2.   |
| 280H             | 640 | IA32_MC0_CTL2              | Module  | See Table 2-2.   |
| 281H             | 641 | IA32_MC1_CTL2              | Module  | See Table 2-2.   |
| 282H             | 642 | IA32_MC2_CTL2              | Core    | See Table 2-2.   |
| 283H             | 643 | IA32_MC3_CTL2              | Module  | See Table 2-2.   |
| 284H             | 644 | IA32_MC4_CTL2              | Package | See Table 2-2.   |
| 285H             | 645 | IA32_MC5_CTL2              | Package | See Table 2-2.   |
| 286H             | 646 | IA32_MC6_CTL2              | Package | See Table 2-2.   |
| 300H             | 768 | MSR_SGXOWNEREPOCH0         | Package | Lower 64 Bit CR_SGXOWNEREPOCH (w)<br>Writes do not update CR_SGXOWNEREPOCH if CPUID.(EAX=12H, ECX=0):EAX.SGX1 is 1 on any thread in the package.                                 |
|                  |     | 63:0                       |         | Lower 64 bits of an 128-bit external entropy value for key derivation of an enclave.   |
| 301H             | 769 | MSR_SGXOWNEREPOCH1         | Package | Upper 64 Bit CR_SGXOWNEREPOCH (w)<br>Writes do not update CR_SGXOWNEREPOCH if CPUID.(EAX=12H, ECX=0):EAX.SGX1 is 1 on any thread in the package.                                 |
|                  |     | 63:0                       |         | Upper 64 bits of an 128-bit external entropy value for key derivation of an enclave.   |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |     | Register Name / Bit Fields    | Scope | Bit Description  |
|------------------|-----|-------------------------------|-------|--|
| Hex              | Dec |                               |       |  |
| 38EH             | 910 | IA32_PERF_GLOBAL_STATUS       | Core  | See Table 2-2. See Section 18.2.4, "Architectural Performance Monitoring Version 4." |
|                  |     | 0                             |       | Ovf_PMC0   |
|                  |     | 1                             |       | Ovf_PMC1   |
|                  |     | 2                             |       | Ovf_PMC2   |
|                  |     | 3                             |       | Ovf_PMC3   |
|                  |     | 31:4                          |       | Reserved   |
|                  |     | 32                            |       | Ovf_FixedCtr0  |
|                  |     | 33                            |       | Ovf_FixedCtr1  |
|                  |     | 34                            |       | Ovf_FixedCtr2  |
|                  |     | 54:35                         |       | Reserved   |
|                  |     | 55                            |       | Trace_ToPA_PMI   |
|                  |     | 57:56                         |       | Reserved   |
|                  |     | 58                            |       | LBR_Frz.   |
|                  |     | 59                            |       | CTR_Frz.   |
|                  |     | 60                            |       | ASCI   |
|                  |     | 61                            |       | Ovf_Uncore   |
|                  |     | 62                            |       | Ovf_BufDSSAVE  |
| 63               |     | CondChgd                      |       |  |
| 390H             | 912 | IA32_PERF_GLOBAL_STATUS_RESET | Core  | See Table 2-2. See Section 18.2.4, "Architectural Performance Monitoring Version 4." |
|                  |     | 0                             |       | Set 1 to clear Ovf_PMC0.   |
|                  |     | 1                             |       | Set 1 to clear Ovf_PMC1.   |
|                  |     | 2                             |       | Set 1 to clear Ovf_PMC2.   |
|                  |     | 3                             |       | Set 1 to clear Ovf_PMC3.   |
|                  |     | 31:4                          |       | Reserved   |
|                  |     | 32                            |       | Set 1 to clear Ovf_FixedCtr0.  |
|                  |     | 33                            |       | Set 1 to clear Ovf_FixedCtr1.  |
|                  |     | 34                            |       | Set 1 to clear Ovf_FixedCtr2.  |
|                  |     | 54:35                         |       | Reserved   |
|                  |     | 55                            |       | Set 1 to clear Trace_ToPA_PMI.   |
|                  |     | 57:56                         |       | Reserved   |
|                  |     | 58                            |       | Set 1 to clear LBR_Frz.  |
|                  |     | 59                            |       | Set 1 to clear CTR_Frz.  |
|                  |     | 60                            |       | Set 1 to clear ASCI.   |
|                  |     | 61                            |       | Set 1 to clear Ovf_Uncore.   |
|                  |     | 62                            |       | Set 1 to clear Ovf_BufDSSAVE.  |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields    | Scope   | Bit Description  |
|------------------|------|-------------------------------|---------|--|
| Hex              | Dec  |                               |         |  |
|                  |      | 63                            |         | Set 1 to clear CondChgd.   |
| 391H             | 913  | IA32_PERF_GLOBAL_STATUS_SET   | Core    | See Table 2-2. See Section 18.2.4, "Architectural Performance Monitoring Version 4."   |
|                  |      | 0                             |         | Set 1 to cause Ovf_PMC0 = 1.   |
|                  |      | 1                             |         | Set 1 to cause Ovf_PMC1 = 1.   |
|                  |      | 2                             |         | Set 1 to cause Ovf_PMC2 = 1.   |
|                  |      | 3                             |         | Set 1 to cause Ovf_PMC3 = 1.   |
|                  |      | 31:4                          |         | Reserved   |
|                  |      | 32                            |         | Set 1 to cause Ovf_FixedCtr0 = 1.  |
|                  |      | 33                            |         | Set 1 to cause Ovf_FixedCtr1 = 1.  |
|                  |      | 34                            |         | Set 1 to cause Ovf_FixedCtr2 = 1.  |
|                  |      | 54:35                         |         | Reserved   |
|                  |      | 55                            |         | Set 1 to cause Trace_ToPA_PMI = 1.   |
|                  |      | 57:56                         |         | Reserved   |
|                  |      | 58                            |         | Set 1 to cause LBR_Frz = 1.  |
|                  |      | 59                            |         | Set 1 to cause CTR_Frz = 1.  |
|                  |      | 60                            |         | Set 1 to cause ASCI = 1.   |
|                  |      | 61                            |         | Set 1 to cause Ovf_Uncore.   |
| 62               |      | Set 1 to cause Ovf_BufDSSAVE. |         |  |
| 63               |      | Reserved                      |         |  |
| 392H             | 914  | IA32_PERF_GLOBAL_INUSE        |         | See Table 2-2.   |
| 3F1H             | 1009 | MSR_PEBS_ENABLE               | Core    | See Table 2-2. See Section 18.6.2.4, "Processor Event Based Sampling (PEBS)."  |
|                  |      | 0                             |         | Enable PEBS trigger and recording for the programmed event (precise or otherwise) on IA32_PMC0. (R/W)  |
| 3F8H             | 1016 | MSR_PKG_C3_RESIDENCY          | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                          |
|                  |      | 63:0                          |         | Package C3 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C3 states. Count at the same frequency as the TSC. |
| 3F9H             | 1017 | MSR_PKG_C6_RESIDENCY          | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                          |
|                  |      | 63:0                          |         | Package C6 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C6 states. Count at the same frequency as the TSC. |



Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 3FCH             | 1020 | MSR_CORE_C3_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |
|                  |      | 63:0                       |         | CORE C3 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C3 states. Count at the same frequency as the TSC.   |
| 406H             | 1030 | IA32_MC1_ADDR              | Module  | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.  |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 419H             | 1049 | IA32_MC6_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 4C3H             | 1219 | IA32_A_PMC2                | Core    | See Table 2-2.   |
| 4C4H             | 1220 | IA32_A_PMC3                | Core    | See Table 2-2.   |
| 4E0H             | 1248 | MSR_SMM_FEATURE_CONTROL    | Package | Enhanced SMM Feature Control (SMM-RW)<br>Reports SMM capability Enhancement. Accessible only while in SMM.   |
|                  |      | 0                          |         | Lock (SMM-RW0)<br>When set to '1' locks this register from further changes.  |
|                  |      | 1                          |         | Reserved   |
|                  |      | 2                          |         | SMM_Code_Chk_En (SMM-RW)<br>This control bit is available only if MSR_SMM_MCA_CAP[58] == 1. When set to '0' (default) none of the logical processors are prevented from executing SMM code outside the ranges defined by the SMRR.<br>When set to '1' any logical processor in the package that attempts to execute SMM code not within the ranges defined by the SMRR will assert an unrecoverable MCE. |
|                  |      | 63:3                       |         | Reserved   |
| 4E2H             | 1250 | MSR_SMM_DELAYED            | Package | SMM Delayed (SMM-RO)<br>Reports the interruptible state of all logical processors in the package. Available only while in SMM and MSR_SMM_MCA_CAP[LONG_FLOW_INDICATION] == 1.  |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | N-1:0                      |         | LOG_PROC_STATE (SMM-RO)<br>Each bit represents a processor core of its state in a long flow of internal operation which delays servicing an interrupt. The corresponding bit will be set at the start of long events such as: Microcode Update Load, C6, WBINVD, Ratio Change, Throttle.<br>The bit is automatically cleared at the end of each long event. The reset value of this field is 0.<br>Only bit positions below N = CPUID.(EAX=0BH, ECX=PKG_LVL):EBX[15:0] can be updated. |
|                  |      | 63:N                       |         | Reserved   |
| 4E3H             | 1251 | MSR_SMM_BLOCKED            | Package | SMM Blocked (SMM-RO)<br>Reports the blocked state of all logical processors in the package. Available only while in SMM.   |
|                  |      | N-1:0                      |         | LOG_PROC_STATE (SMM-RO)<br>Each bit represents a processor core of its blocked state to service an SMI. The corresponding bit will be set if the logical processor is in one of the following states: Wait For SIPI or SENTER Sleep.<br>The reset value of this field is OFFFH.<br>Only bit positions below N = CPUID.(EAX=0BH, ECX=PKG_LVL):EBX[15:0] can be updated.   |
|                  |      | 63:N                       |         | Reserved   |
| 500H             | 1280 | IA32_SGX_SVN_STATUS        | Core    | Status and SVN Threshold of SGX Support for ACM (RO)   |
|                  |      | 0                          |         | Lock<br>See Section 4.1.1.3, "Interactions with Authenticated Code Modules (ACMs)".  |
|                  |      | 15:1                       |         | Reserved   |
|                  |      | 23:16                      |         | SGX_SVN_SINIT<br>See Section 4.1.1.3, "Interactions with Authenticated Code Modules (ACMs)".   |
|                  |      | 63:24                      |         | Reserved   |
| 560H             | 1376 | IA32_RTIT_OUTPUT_BASE      | Core    | Trace Output Base Register (R/W)<br>See Table 2-2.   |
| 561H             | 1377 | IA32_RTIT_OUTPUT_MASK_PTRS | Core    | Trace Output Mask Pointers Register (R/W)<br>See Table 2-2.  |
| 570H             | 1392 | IA32_RTIT_CTL              | Core    | Trace Control Register (R/W)   |
|                  |      | 0                          |         | TraceEn  |
|                  |      | 1                          |         | CYCEn  |
|                  |      | 2                          |         | OS   |
|                  |      | 3                          |         | User   |
|                  |      | 6:4                        |         | Reserved, must be zero.  |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description                                     |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
|                  |      | 7                          |       | CR3 filter  |
|                  |      | 8                          |       | ToPA<br>Writing 0 will #GP if also setting TraceEn. |
|                  |      | 9                          |       | MTCEn   |
|                  |      | 10                         |       | TSCEn   |
|                  |      | 11                         |       | DisRETC   |
|                  |      | 12                         |       | Reserved, must be zero.                             |
|                  |      | 13                         |       | BranchEn  |
|                  |      | 17:14                      |       | MTCFreq   |
|                  |      | 18                         |       | Reserved, must be zero.                             |
|                  |      | 22:19                      |       | CYCThresh   |
|                  |      | 23                         |       | Reserved, must be zero.                             |
|                  |      | 27:24                      |       | PSBFreq   |
|                  |      | 31:28                      |       | Reserved, must be zero.                             |
|                  |      | 35:32                      |       | ADDRO_CFG   |
|                  |      | 39:36                      |       | ADDR1_CFG   |
|                  |      | 63:40                      |       | Reserved, must be zero.                             |
| 571H             | 1393 | IA32_RTIT_STATUS           | Core  | Tracing Status Register (R/W)                       |
|                  |      | 0                          |       | FilterEn<br>Writes ignored.                         |
|                  |      | 1                          |       | ContexEn<br>Writes ignored.                         |
|                  |      | 2                          |       | TriggerEn<br>Writes ignored.                        |
|                  |      | 3                          |       | Reserved  |
|                  |      | 4                          |       | Error (R/W)   |
|                  |      | 5                          |       | Stopped   |
|                  |      | 31:6                       |       | Reserved, must be zero.                             |
|                  |      | 48:32                      |       | PacketByteCnt                                       |
|                  |      | 63:49                      |       | Reserved, must be zero.                             |
| 572H             | 1394 | IA32_RTIT_CR3_MATCH        | Core  | Trace Filter CR3 Match Register (R/W)               |
|                  |      | 4:0                        |       | Reserved  |
|                  |      | 63:5                       |       | CR3[63:5] value to match.                           |
| 580H             | 1408 | IA32_RTIT_ADDRO_A          | Core  | Region 0 Start Address (R/W)                        |
|                  |      | 63:0                       |       | See Table 2-2.                                      |
| 581H             | 1409 | IA32_RTIT_ADDRO_B          | Core  | Region 0 End Address (R/W)                          |
|                  |      | 63:0                       |       | See Table 2-2.                                      |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 582H             | 1410 | IA32_RTIT_ADDR1_A          | Core    | Region 1 Start Address (R/W)  |
|                  |      | 63:0                       |         | See Table 2-2.  |
| 583H             | 1411 | IA32_RTIT_ADDR1_B          | Core    | Region 1 End Address (R/W)  |
|                  |      | 63:0                       |         | See Table 2-2.  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers used in RAPL Interfaces (R/O)<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 3:0                        |         | Power Units<br>Power related information (in Watts) is in unit of $1W/2^{PU}$ ; where PU is an unsigned integer represented by bits 3:0. Default value is 1000b, indicating power unit is in 3.9 milliWatts increment.          |
|                  |      | 7:4                        |         | Reserved  |
|                  |      | 12:8                       |         | Energy Status Units<br>Energy related information (in Joules) is in unit of $1Joule/2^{ESU}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 01110b, indicating energy unit is in 61 microJoules. |
|                  |      | 15:13                      |         | Reserved  |
|                  |      | 19:16                      |         | Time Unit<br>Time related information (in seconds) is in unit of $1S/2^{TU}$ ; where TU is an unsigned integer represented by bits 19:16. Default value is 1010b, indicating power unit is in 0.977 millisecond.                |
|                  |      | 63:20                      |         | Reserved  |
| 60AH             | 1546 | MSR_PKGC3_IRTL             | Package | Package C3 Interrupt Response Limit (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C3 state.  |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. See Table 2-19 for supported time unit encodings.  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-sate management.  |
|                  |      | 63:16                      |         | Reserved  |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 60BH             | 1547 | MSR_PKG_C2_IRTL1           | Package | Package C6/C7S Interrupt Response Limit 1 (R/W)<br>This MSR defines the interrupt response time limit used by the processor to manage a transition to a package C6 or C7S state.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states. |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C6 or C7S state.   |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. See Table 2-19 for supported time unit encodings.  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |
|                  |      | 63:16                      |         | Reserved  |
| 60CH             | 1548 | MSR_PKG_C2_IRTL2           | Package | Package C7 Interrupt Response Limit 2 (R/W)<br>This MSR defines the interrupt response time limit used by the processor to manage a transition to a package C7 state.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.            |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C7 state.  |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. See Table 2-19 for supported time unit encodings.  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |
|                  |      | 63:16                      |         | Reserved  |
| 60DH             | 1549 | MSR_PKG_C2_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:0                       |         | Package C2 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C2 states. Count at the same frequency as the TSC.  |
| 610H             | 1552 | MSR_PKG_POWER_LIMIT        | Package | PKG RAPL Power Limit Control (R/W)<br>See Section 14.9.3, "Package RAPL Domain."  |
| 611H             | 1553 | MSR_PKG_ENERGY_STATUS      | Package | PKG Energy Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 613H             | 1555 | MSR_PKG_PERF_STATUS        | Package | PKG Perf Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 614H             | 1556 | MSR_PKG_POWER_INFO         | Package | PKG RAPL Parameters (R/W)   |
|                  |      | 14:0                       |         | Thermal Spec Power (R/W)<br>See Section 14.9.3, "Package RAPL Domain."  |
|                  |      | 15                         |         | Reserved  |
|                  |      | 30:16                      |         | Minimum Power (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
|                  |      | 31                         |         | Reserved  |
|                  |      | 46:32                      |         | Maximum Power (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
|                  |      | 47                         |         | Reserved  |
|                  |      | 54:48                      |         | Maximum Time Window (R/W)<br>Specified by $2^Y * (1.0 + Z/4.0) * \text{Time\_Unit}$ , where "Y" is the unsigned integer value represented by bits 52:48, "Z" is an unsigned integer represented by bits 54:53. "Time_Unit" is specified by the "Time Units" field of MSR_RAPL_POWER_UNIT. |
|                  |      | 63:55                      |         | Reserved  |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 632H             | 1586 |                            | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |
|                  |      | 63:0                       |         | Package C10 Residency Counter (R/O)<br>Value since last reset that the entire SOC is in an S0i3 state. Count at the same frequency as the TSC.  |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields   | Scope   | Bit Description  |
|------------------|------|--|---------|--|
| Hex              | Dec  |  |         |  |
| 639H             | 1593 | MSR_PP0_ENERGY_STATUS  | Package | PP0 Energy Status (R/O)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."   |
| 641H             | 1601 | MSR_PP1_ENERGY_STATUS  | Package | PP1 Energy Status (R/O)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."   |
| 64CH             | 1612 | MSR_TURBO_ACTIVATION_RATIO   | Package | ConfigTDP Control (R/W)  |
|                  |      | 7:0  |         | MAX_NON_TURBO_RATIO (R/W/L)<br>System BIOS can program this field.   |
|                  |      | 30:8   |         | Reserved   |
|                  |      | 31   |         | TURBO_ACTIVATION_RATIO_Lock (R/W/L)<br>When this bit is set, the content of this register is locked until a reset.   |
|                  |      | 63:32  |         | Reserved   |
| 64FH             | 1615 | MSR_CORE_PERF_LIMIT_REASONS  | Package | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.)  |
|                  |      | 0  |         | PROCHOT Status (R0)<br>When set, processor core frequency is reduced below the operating system request due to assertion of external PROCHOT.              |
|                  |      | 1  |         | Thermal Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal event.   |
|                  |      | 2  |         | Package-Level Power Limiting PL1 Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL1. |
|                  |      | 3  |         | Package-Level PL2 Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL2. |
|                  |      | 8:4  |         | Reserved   |
|                  |      | 9  |         | Core Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to domain-level power limiting.                   |
|                  |      | 10   |         | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.         |
| 11               |      | Max Turbo Limit Status (R0)<br>When set, frequency is reduced below the operating system request due to multi-core turbo limits. |         |  |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 12                         |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption).                                 |
|                  |     | 13                         |       | Turbo Transition Attenuation Status (R0)<br>When set, frequency is reduced below the operating system request due to Turbo transition attenuation. This prevents performance degradation due to frequent operating ratio changes.   |
|                  |     | 14                         |       | Maximum Efficiency Frequency Status (R0)<br>When set, frequency is reduced below the maximum efficiency frequency.  |
|                  |     | 15                         |       | Reserved  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared. This log bit will remain set until cleared by software writing 0.   |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared. This log bit will remain set until cleared by software writing 0.   |
|                  |     | 18                         |       | Package-Level PL1 Power Limiting Log<br>When set, indicates that the Package Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared. This log bit will remain set until cleared by software writing 0. |
|                  |     | 19                         |       | Package-Level PL2 Power Limiting Log<br>When set, indicates that the Package Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared. This log bit will remain set until cleared by software writing 0. |
|                  |     | 24:20                      |       | Reserved  |
|                  |     | 25                         |       | Core Power Limiting Log<br>When set, indicates that the Core Power Limiting Status bit has asserted since the log bit was last cleared. This log bit will remain set until cleared by software writing 0.                           |



Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
|                  |      | 26                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 27                         |       | Max Turbo Limit Log<br>When set, indicates that the Max Turbo Limit Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 28                         |       | Electrical Design Point Log<br>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 29                         |       | Turbo Transition Attenuation Log<br>When set, indicates that the Turbo Transition Attenuation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 30                         |       | Maximum Efficiency Frequency Log<br>When set, indicates that the Maximum Efficiency Frequency Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 63:31                      |       | Reserved  |
| 680H             | 1664 | MSR_LASTBRANCH_0_FROM_IP   | Core  | Last Branch Record 0 From IP (R/W)<br>One of 32 pairs of last branch record registers on the last branch record stack. The From_IP part of the stack contains pointers to the source instruction . See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.6 and record format in Section 17.4.8.1.</li> </ul> |
|                  |      | 0:47                       |       | From Linear Address (R/W)   |
|                  |      | 62:48                      |       | Signed extension of bits 47:0.  |
|                  |      | 63                         |       | Mispred   |
| 681H             | 1665 | MSR_LASTBRANCH_1_FROM_IP   | Core  | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 682H             | 1666 | MSR_LASTBRANCH_2_FROM_IP   | Core  | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 683H             | 1667 | MSR_LASTBRANCH_3_FROM_IP   | Core  | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| 684H             | 1668 | MSR_LASTBRANCH_4_FROM_IP   | Core  | Last Branch Record 4 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 685H             | 1669 | MSR_LASTBRANCH_5_FROM_IP   | Core  | Last Branch Record 5 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 686H             | 1670 | MSR_LASTBRANCH_6_FROM_IP   | Core  | Last Branch Record 6 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 687H             | 1671 | MSR_LASTBRANCH_7_FROM_IP   | Core  | Last Branch Record 7 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 688H             | 1672 | MSR_LASTBRANCH_8_FROM_IP   | Core  | Last Branch Record 8 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 689H             | 1673 | MSR_LASTBRANCH_9_FROM_IP   | Core  | Last Branch Record 9 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68AH             | 1674 | MSR_LASTBRANCH_10_FROM_IP  | Core  | Last Branch Record 10 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 68BH             | 1675 | MSR_LASTBRANCH_11_FROM_IP  | Core  | Last Branch Record 11 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 68CH             | 1676 | MSR_LASTBRANCH_12_FROM_IP  | Core  | Last Branch Record 12 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 68DH             | 1677 | MSR_LASTBRANCH_13_FROM_IP  | Core  | Last Branch Record 13 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 68EH             | 1678 | MSR_LASTBRANCH_14_FROM_IP  | Core  | Last Branch Record 14 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 68FH             | 1679 | MSR_LASTBRANCH_15_FROM_IP  | Core  | Last Branch Record 15 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 690H             | 1680 | MSR_LASTBRANCH_16_FROM_IP  | Core  | Last Branch Record 16 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 691H             | 1681 | MSR_LASTBRANCH_17_FROM_IP  | Core  | Last Branch Record 17 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 692H             | 1682 | MSR_LASTBRANCH_18_FROM_IP  | Core  | Last Branch Record 18 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 693H             | 1683 | MSR_LASTBRANCH_19_FROM_IP  | Core  | Last Branch Record 19 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 694H             | 1684 | MSR_LASTBRANCH_20_FROM_IP  | Core  | Last Branch Record 20 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 695H             | 1685 | MSR_LASTBRANCH_21_FROM_IP  | Core  | Last Branch Record 21 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |
| 696H             | 1686 | MSR_LASTBRANCH_22_FROM_IP  | Core  | Last Branch Record 22 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP. |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|------|----------------------------|-------|--|
| Hex              | Dec  |                            |       |  |
| 697H             | 1687 | MSR_LASTBRANCH_23_FROM_IP  | Core  | Last Branch Record 23 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 698H             | 1688 | MSR_LASTBRANCH_24_FROM_IP  | Core  | Last Branch Record 24 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 699H             | 1689 | MSR_LASTBRANCH_25_FROM_IP  | Core  | Last Branch Record 25 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69AH             | 1690 | MSR_LASTBRANCH_26_FROM_IP  | Core  | Last Branch Record 26 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69BH             | 1691 | MSR_LASTBRANCH_27_FROM_IP  | Core  | Last Branch Record 27 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69CH             | 1692 | MSR_LASTBRANCH_28_FROM_IP  | Core  | Last Branch Record 28 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69DH             | 1693 | MSR_LASTBRANCH_29_FROM_IP  | Core  | Last Branch Record 29 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69EH             | 1694 | MSR_LASTBRANCH_30_FROM_IP  | Core  | Last Branch Record 30 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69FH             | 1695 | MSR_LASTBRANCH_31_FROM_IP  | Core  | Last Branch Record 31 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 6C0H             | 1728 | MSR_LASTBRANCH_0_TO_IP     | Core  | Last Branch Record 0 To IP (R/W)<br>One of 32 pairs of last branch record registers on the last branch record stack. The To_IP part of the stack contains pointers to the Destination instruction and elapsed cycles from last LBR update. See Section 17.6. |
|                  |      | 0:47                       |       | Target Linear Address (R/W)  |
|                  |      | 63:48                      |       | Elapsed cycles from last update to the LBR.  |
| 6C1H             | 1729 | MSR_LASTBRANCH_1_TO_IP     | Core  | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C2H             | 1730 | MSR_LASTBRANCH_2_TO_IP     | Core  | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C3H             | 1731 | MSR_LASTBRANCH_3_TO_IP     | Core  | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C4H             | 1732 | MSR_LASTBRANCH_4_TO_IP     | Core  | Last Branch Record 4 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C5H             | 1733 | MSR_LASTBRANCH_5_TO_IP     | Core  | Last Branch Record 5 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C6H             | 1734 | MSR_LASTBRANCH_6_TO_IP     | Core  | Last Branch Record 6 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C7H             | 1735 | MSR_LASTBRANCH_7_TO_IP     | Core  | Last Branch Record 7 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| 6C8H             | 1736 | MSR_LASTBRANCH_8_TO_IP     | Core  | Last Branch Record 8 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6C9H             | 1737 | MSR_LASTBRANCH_9_TO_IP     | Core  | Last Branch Record 9 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6CAH             | 1738 | MSR_LASTBRANCH_10_TO_IP    | Core  | Last Branch Record 10 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CBH             | 1739 | MSR_LASTBRANCH_11_TO_IP    | Core  | Last Branch Record 11 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CCH             | 1740 | MSR_LASTBRANCH_12_TO_IP    | Core  | Last Branch Record 12 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CDH             | 1741 | MSR_LASTBRANCH_13_TO_IP    | Core  | Last Branch Record 13 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CEH             | 1742 | MSR_LASTBRANCH_14_TO_IP    | Core  | Last Branch Record 14 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CFH             | 1743 | MSR_LASTBRANCH_15_TO_IP    | Core  | Last Branch Record 15 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DOH             | 1744 | MSR_LASTBRANCH_16_TO_IP    | Core  | Last Branch Record 16 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D1H             | 1745 | MSR_LASTBRANCH_17_TO_IP    | Core  | Last Branch Record 17 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D2H             | 1746 | MSR_LASTBRANCH_18_TO_IP    | Core  | Last Branch Record 18 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D3H             | 1747 | MSR_LASTBRANCH_19_TO_IP    | Core  | Last Branch Record 19 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D4H             | 1748 | MSR_LASTBRANCH_20_TO_IP    | Core  | Last Branch Record 20 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D5H             | 1749 | MSR_LASTBRANCH_21_TO_IP    | Core  | Last Branch Record 21 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D6H             | 1750 | MSR_LASTBRANCH_22_TO_IP    | Core  | Last Branch Record 22 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D7H             | 1751 | MSR_LASTBRANCH_23_TO_IP    | Core  | Last Branch Record 23 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D8H             | 1752 | MSR_LASTBRANCH_24_TO_IP    | Core  | Last Branch Record 24 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6D9H             | 1753 | MSR_LASTBRANCH_25_TO_IP    | Core  | Last Branch Record 25 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DAH             | 1754 | MSR_LASTBRANCH_26_TO_IP    | Core  | Last Branch Record 26 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |

Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| 6DBH             | 1755 | MSR_LASTBRANCH_27_TO_IP    | Core  | Last Branch Record 27 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DCH             | 1756 | MSR_LASTBRANCH_28_TO_IP    | Core  | Last Branch Record 28 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DDH             | 1757 | MSR_LASTBRANCH_29_TO_IP    | Core  | Last Branch Record 29 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DEH             | 1758 | MSR_LASTBRANCH_30_TO_IP    | Core  | Last Branch Record 30 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6DFH             | 1759 | MSR_LASTBRANCH_31_TO_IP    | Core  | Last Branch Record 31 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 802H             | 2050 | IA32_X2APIC_APICID         | Core  | x2APIC ID register (R/O)  |
| 803H             | 2051 | IA32_X2APIC_VERSION        | Core  | x2APIC Version register (R/O)   |
| 808H             | 2056 | IA32_X2APIC_TPR            | Core  | x2APIC Task Priority register (R/W)   |
| 80AH             | 2058 | IA32_X2APIC_PPR            | Core  | x2APIC Processor Priority register (R/O)  |
| 80BH             | 2059 | IA32_X2APIC_EOI            | Core  | x2APIC EOI register (W/O)   |
| 80DH             | 2061 | IA32_X2APIC_LDR            | Core  | x2APIC Logical Destination register (R/O)                                       |
| 80FH             | 2063 | IA32_X2APIC_SIVR           | Core  | x2APIC Spurious Interrupt Vector register (R/W)                                 |
| 810H             | 2064 | IA32_X2APIC_ISR0           | Core  | x2APIC In-Service register bits [31:0] (R/O)                                    |
| 811H             | 2065 | IA32_X2APIC_ISR1           | Core  | x2APIC In-Service register bits [63:32] (R/O)                                   |
| 812H             | 2066 | IA32_X2APIC_ISR2           | Core  | x2APIC In-Service register bits [95:64] (R/O)                                   |
| 813H             | 2067 | IA32_X2APIC_ISR3           | Core  | x2APIC In-Service register bits [127:96] (R/O)                                  |
| 814H             | 2068 | IA32_X2APIC_ISR4           | Core  | x2APIC In-Service register bits [159:128] (R/O)                                 |
| 815H             | 2069 | IA32_X2APIC_ISR5           | Core  | x2APIC In-Service register bits [191:160] (R/O)                                 |
| 816H             | 2070 | IA32_X2APIC_ISR6           | Core  | x2APIC In-Service register bits [223:192] (R/O)                                 |
| 817H             | 2071 | IA32_X2APIC_ISR7           | Core  | x2APIC In-Service register bits [255:224] (R/O)                                 |
| 818H             | 2072 | IA32_X2APIC_TMR0           | Core  | x2APIC Trigger Mode register bits [31:0] (R/O)                                  |
| 819H             | 2073 | IA32_X2APIC_TMR1           | Core  | x2APIC Trigger Mode register bits [63:32] (R/O)                                 |
| 81AH             | 2074 | IA32_X2APIC_TMR2           | Core  | x2APIC Trigger Mode register bits [95:64] (R/O)                                 |
| 81BH             | 2075 | IA32_X2APIC_TMR3           | Core  | x2APIC Trigger Mode register bits [127:96] (R/O)                                |
| 81CH             | 2076 | IA32_X2APIC_TMR4           | Core  | x2APIC Trigger Mode register bits [159:128] (R/O)                               |
| 81DH             | 2077 | IA32_X2APIC_TMR5           | Core  | x2APIC Trigger Mode register bits [191:160] (R/O)                               |
| 81EH             | 2078 | IA32_X2APIC_TMR6           | Core  | x2APIC Trigger Mode register bits [223:192] (R/O)                               |
| 81FH             | 2079 | IA32_X2APIC_TMR7           | Core  | x2APIC Trigger Mode register bits [255:224] (R/O)                               |
| 820H             | 2080 | IA32_X2APIC_IRR0           | Core  | x2APIC Interrupt Request register bits [31:0] (R/O)                             |
| 821H             | 2081 | IA32_X2APIC_IRR1           | Core  | x2APIC Interrupt Request register bits [63:32] (R/O)                            |
| 822H             | 2082 | IA32_X2APIC_IRR2           | Core  | x2APIC Interrupt Request register bits [95:64] (R/O)                            |
| 823H             | 2083 | IA32_X2APIC_IRR3           | Core  | x2APIC Interrupt Request register bits [127:96] (R/O)                           |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| 824H             | 2084 | IA32_X2APIC_IRR4           | Core   | x2APIC Interrupt Request register bits [159:128] (R/O)                                     |
| 825H             | 2085 | IA32_X2APIC_IRR5           | Core   | x2APIC Interrupt Request register bits [191:160] (R/O)                                     |
| 826H             | 2086 | IA32_X2APIC_IRR6           | Core   | x2APIC Interrupt Request register bits [223:192] (R/O)                                     |
| 827H             | 2087 | IA32_X2APIC_IRR7           | Core   | x2APIC Interrupt Request register bits [255:224] (R/O)                                     |
| 828H             | 2088 | IA32_X2APIC_ESR            | Core   | x2APIC Error Status register (R/W)   |
| 82FH             | 2095 | IA32_X2APIC_LVT_CMCI       | Core   | x2APIC LVT Corrected Machine Check Interrupt register (R/W)                                |
| 830H             | 2096 | IA32_X2APIC_ICR            | Core   | x2APIC Interrupt Command register (R/W)  |
| 832H             | 2098 | IA32_X2APIC_LVT_TIMER      | Core   | x2APIC LVT Timer Interrupt register (R/W)  |
| 833H             | 2099 | IA32_X2APIC_LVT_THERMAL    | Core   | x2APIC LVT Thermal Sensor Interrupt register (R/W)   |
| 834H             | 2100 | IA32_X2APIC_LVT_PMI        | Core   | x2APIC LVT Performance Monitor register (R/W)  |
| 835H             | 2101 | IA32_X2APIC_LVT_LINT0      | Core   | x2APIC LVT LINT0 register (R/W)  |
| 836H             | 2102 | IA32_X2APIC_LVT_LINT1      | Core   | x2APIC LVT LINT1 register (R/W)  |
| 837H             | 2103 | IA32_X2APIC_LVT_ERROR      | Core   | x2APIC LVT Error register (R/W)  |
| 838H             | 2104 | IA32_X2APIC_INIT_COUNT     | Core   | x2APIC Initial Count register (R/W)  |
| 839H             | 2105 | IA32_X2APIC_CUR_COUNT      | Core   | x2APIC Current Count register (R/O)  |
| 83EH             | 2110 | IA32_X2APIC_DIV_CONF       | Core   | x2APIC Divide Configuration register (R/W)   |
| 83FH             | 2111 | IA32_X2APIC_SELF_IPI       | Core   | x2APIC Self IPI register (W/O)   |
| C8FH             | 3215 | IA32_PQR_ASSOC             | Core   | Resource Association Register (R/W)  |
|                  |      | 31:0                       |        | Reserved   |
|                  |      | 33:32                      |        | COS (R/W)  |
|                  |      | 63:34                      |        | Reserved   |
| D10H             | 3344 | IA32_L2_QOS_MASK_0         | Module | L2 Class Of Service Mask - COS 0 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=0. |
|                  |      | 0:7                        |        | CBM: Bit vector of available L2 ways for COS 0 enforcement.                                |
|                  |      | 63:8                       |        | Reserved   |
| D11H             | 3345 | IA32_L2_QOS_MASK_1         | Module | L2 Class Of Service Mask - COS 1 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=1. |
|                  |      | 0:7                        |        | CBM: Bit vector of available L2 ways for COS 0 enforcement.                                |
|                  |      | 63:8                       |        | Reserved   |
| D12H             | 3346 | IA32_L2_QOS_MASK_2         | Module | L2 Class Of Service Mask - COS 2 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=2. |
|                  |      | 0:7                        |        | CBM: Bit vector of available L2 ways for COS 0 enforcement.                                |

**Table 2-12. MSRs in Intel Atom Processors Based on the Goldmont Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 63:8                       |         | Reserved   |
| D13H             | 3347 | IA32_L2_QOS_MASK_3         | Package | L2 Class Of Service Mask - COS 3 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=3. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L2 ways for COS 3 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| D90H             | 3472 | IA32_BNDCFGS               | Core    | See Table 2-2.   |
| DA0H             | 3488 | IA32_XSS                   | Core    | See Table 2-2.   |

See Table 2-6, and Table 2-12 for MSR definitions applicable to processors with CPUID signature 06\_5CH.

## 2.6 MSRS IN INTEL ATOM PROCESSORS BASED ON GOLDMONT PLUS MICROARCHITECTURE

Intel Atom processors based on the Goldmont Plus microarchitecture support MSRs listed in Table 2-6, Table 2-12 and Table 2-13. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_7AH; see Table 2-1. For an MSR listed in Table 2-13 that also appears in the model-specific tables of prior generations, Table 2-13 supercede prior generation tables.

In the Goldmont Plus microarchitecture, the scope column indicates the following: “Core” means each processor core has a separate MSR, or a bit field not shared with another processor core. “Module” means the MSR or the bit field is shared by a pair of processor cores in the physical package. “Package” means all processor cores in the physical package share the same MSR or bit interface.

**Table 2-13. MSRs in Intel Atom Processors Based on the Goldmont Plus Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Core  | Control Features in Intel 64Processor (R/W)<br>See Table 2-2.   |
|                  |     | 0                          |       | Lock (R/WL)   |
|                  |     | 1                          |       | Enable VMX inside SMX operation (R/WL)  |
|                  |     | 2                          |       | Enable VMX outside SMX operation (R/WL)   |
|                  |     | 14:8                       |       | SENTER local functions enables (R/WL)   |
|                  |     | 15                         |       | SENTER global functions enable (R/WL)   |
|                  |     | 17                         |       | SGX Launch Control Enable (R/WL)<br>This bit must be set to enable runtime reconfiguration of SGX Launch Control via IA32_SGXLEPUBKEYHASHn MSR.<br>Valid if CPUID.(EAX=07H, ECX=0H): ECX[30] = 1. |
|                  |     | 18                         |       | SGX global functions enable (R/WL)  |
|                  |     | 63:19                      |       | Reserved  |
| 8CH              | 140 | IA32_SGXLEPUBKEYHASH0      | Core  | See Table 2-2.  |

**Table 2-13. MSRs in Intel Atom Processors Based on the Goldmont Plus Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| 8DH              | 141  | IA32_SGXLEPUBKEYHASH1      | Core  | See Table 2-2.  |
| 8EH              | 142  | IA32_SGXLEPUBKEYHASH2      | Core  | See Table 2-2.  |
| 8FH              | 143  | IA32_SGXLEPUBKEYHASH3      | Core  | See Table 2-2.  |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Core  | (R/W) See Table 2-2. See Section 18.6.2.4, "Processor Event Based Sampling (PEBS)."             |
|                  |      | 0                          |       | Enable PEBS trigger and recording for the programmed event (precise or otherwise) on IA32_PMC0. |
|                  |      | 1                          |       | Enable PEBS trigger and recording for the programmed event (precise or otherwise) on IA32_PMC1. |
|                  |      | 2                          |       | Enable PEBS trigger and recording for the programmed event (precise or otherwise) on IA32_PMC2. |
|                  |      | 3                          |       | Enable PEBS trigger and recording for the programmed event (precise or otherwise) on IA32_PMC3. |
|                  |      | 31:4                       |       | Reserved  |
|                  |      | 32                         |       | Enable PEBS trigger and recording for IA32_FIXED_CTR0.  |
|                  |      | 33                         |       | Enable PEBS trigger and recording for IA32_FIXED_CTR1.  |
|                  |      | 34                         |       | Enable PEBS trigger and recording for IA32_FIXED_CTR2.  |
|                  |      | 63:35                      |       | Reserved  |
| 570H             | 1392 | IA32_RTIT_CTL              | Core  | Trace Control Register (R/W)  |
|                  |      | 0                          |       | TraceEn   |
|                  |      | 1                          |       | CYCEn   |
|                  |      | 2                          |       | OS  |
|                  |      | 3                          |       | User  |
|                  |      | 4                          |       | PwrEvtEn  |
|                  |      | 5                          |       | FUPonPTW  |
|                  |      | 6                          |       | FabricEn  |
|                  |      | 7                          |       | CR3 filter  |
|                  |      | 8                          |       | ToPA<br>Writing 0 will #GP if also setting TraceEn.   |
|                  |      | 9                          |       | MTCEn   |
|                  |      | 10                         |       | TSCEn   |
|                  |      | 11                         |       | DisRETC   |
|                  |      | 12                         |       | PTWEn   |
| 13               |      | BranchEn                   |       |   |



Table 2-13. MSRs in Intel Atom Processors Based on the Goldmont Plus Microarchitecture (Contd.)

| Register Address  |                   | Register Name / Bit Fields | Scope | Bit Description  |
|-------------------|-------------------|----------------------------|-------|--|
| Hex               | Dec               |                            |       |  |
|                   |                   | 17:14                      |       | MTCFreq  |
|                   |                   | 18                         |       | Reserved, must be zero.  |
|                   |                   | 22:19                      |       | CYCThresh  |
|                   |                   | 23                         |       | Reserved, must be zero.  |
|                   |                   | 27:24                      |       | PSBFreq  |
|                   |                   | 31:28                      |       | Reserved, must be zero.  |
|                   |                   | 35:32                      |       | ADDRO_CFG  |
|                   |                   | 39:36                      |       | ADDR1_CFG  |
|                   |                   | 63:40                      |       | Reserved, must be zero.  |
| 680H              | 1664              | MSR_LASTBRANCH_0_FROM_IP   | Core  | Last Branch Record 0 From IP (R/W)<br>One of the three MSRs that make up the first entry of the 32-entry LBR stack. The From_IP part of the stack contains pointers to the source instruction. See also: <ul style="list-style-type: none"> <li>Last Branch Record Stack TOS at 1C9H.</li> <li>Section 17.7, "Last Branch, Call Stack, Interrupt, and Exception Recording for Processors based on Goldmont Plus Microarchitecture."</li> </ul> |
| 681H<br>-<br>69FH | 1665<br>-<br>1695 | MSR_LASTBRANCH_i_FROM_IP   | Core  | Last Branch Record <i>i</i> From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP; <i>i</i> = 1-31.   |
| 6C0H              | 1728              | MSR_LASTBRANCH_0_TO_IP     | Core  | Last Branch Record 0 To IP (R/W)<br>One of the three MSRs that make up the first entry of the 32-entry LBR stack. The To_IP part of the stack contains pointers to the Destination instruction. See also: <ul style="list-style-type: none"> <li>Section 17.7, "Last Branch, Call Stack, Interrupt, and Exception Recording for Processors based on Goldmont Plus Microarchitecture."</li> </ul>   |
| 6C1H<br>-<br>6DFH | 1729<br>-<br>1759 | MSR_LASTBRANCH_i_TO_IP     | Core  | Last Branch Record <i>i</i> To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP; <i>i</i> = 1-31.   |
| DC0H              | 3520              | MSR_LASTBRANCH_INFO_0      | Core  | Last Branch Record 0 Additional Information (R/W)<br>One of the three MSRs that make up the first entry of the 32-entry LBR stack. This part of the stack contains flag and elapsed cycle information. See also: <ul style="list-style-type: none"> <li>Last Branch Record Stack TOS at 1C9H.</li> <li>Section 17.9.1, "LBR Stack."</li> </ul>   |
| DC1H              | 3521              | MSR_LASTBRANCH_INFO_1      | Core  | Last Branch Record 1 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.   |
| DC2H              | 3522              | MSR_LASTBRANCH_INFO_2      | Core  | Last Branch Record 2 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.   |
| DC3H              | 3523              | MSR_LASTBRANCH_INFO_3      | Core  | Last Branch Record 3 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.   |

Table 2-13. MSRs in Intel Atom Processors Based on the Goldmont Plus Microarchitecture (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| DC4H             | 3524 | MSR_LASTBRANCH_INFO_4      | Core  | Last Branch Record 4 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DC5H             | 3525 | MSR_LASTBRANCH_INFO_5      | Core  | Last Branch Record 5 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DC6H             | 3526 | MSR_LASTBRANCH_INFO_6      | Core  | Last Branch Record 6 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DC7H             | 3527 | MSR_LASTBRANCH_INFO_7      | Core  | Last Branch Record 7 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DC8H             | 3528 | MSR_LASTBRANCH_INFO_8      | Core  | Last Branch Record 8 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DC9H             | 3529 | MSR_LASTBRANCH_INFO_9      | Core  | Last Branch Record 9 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0.  |
| DCAH             | 3530 | MSR_LASTBRANCH_INFO_10     | Core  | Last Branch Record 10 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DCBH             | 3531 | MSR_LASTBRANCH_INFO_11     | Core  | Last Branch Record 11 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DCCH             | 3532 | MSR_LASTBRANCH_INFO_12     | Core  | Last Branch Record 12 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DCDH             | 3533 | MSR_LASTBRANCH_INFO_13     | Core  | Last Branch Record 13 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DCEH             | 3534 | MSR_LASTBRANCH_INFO_14     | Core  | Last Branch Record 14 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DCFH             | 3535 | MSR_LASTBRANCH_INFO_15     | Core  | Last Branch Record 15 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD0H             | 3536 | MSR_LASTBRANCH_INFO_16     | Core  | Last Branch Record 16 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD1H             | 3537 | MSR_LASTBRANCH_INFO_17     | Core  | Last Branch Record 17 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD2H             | 3538 | MSR_LASTBRANCH_INFO_18     | Core  | Last Branch Record 18 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD3H             | 3539 | MSR_LASTBRANCH_INFO_19     | Core  | Last Branch Record 19 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD4H             | 3520 | MSR_LASTBRANCH_INFO_20     | Core  | Last Branch Record 20 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD5H             | 3521 | MSR_LASTBRANCH_INFO_21     | Core  | Last Branch Record 21 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD6H             | 3522 | MSR_LASTBRANCH_INFO_22     | Core  | Last Branch Record 22 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |

**Table 2-13. MSRs in Intel Atom Processors Based on the Goldmont Plus Microarchitecture (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|------|----------------------------|-------|---|
| Hex              | Dec  |                            |       |   |
| DD7H             | 3523 | MSR_LASTBRANCH_INFO_23     | Core  | Last Branch Record 23 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD8H             | 3524 | MSR_LASTBRANCH_INFO_24     | Core  | Last Branch Record 24 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DD9H             | 3525 | MSR_LASTBRANCH_INFO_25     | Core  | Last Branch Record 25 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDAH             | 3526 | MSR_LASTBRANCH_INFO_26     | Core  | Last Branch Record 26 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDBH             | 3527 | MSR_LASTBRANCH_INFO_27     | Core  | Last Branch Record 27 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDCH             | 3528 | MSR_LASTBRANCH_INFO_28     | Core  | Last Branch Record 28 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDDH             | 3529 | MSR_LASTBRANCH_INFO_29     | Core  | Last Branch Record 29 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDEH             | 3530 | MSR_LASTBRANCH_INFO_30     | Core  | Last Branch Record 30 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |
| DDFH             | 3531 | MSR_LASTBRANCH_INFO_31     | Core  | Last Branch Record 31 Additional Information (R/W)<br>See description of MSR_LASTBRANCH_INFO_0. |

See Table 2-6, Table 2-12 and Table 2-13 for MSR definitions applicable to processors with CPUID signature 06\_7AH.

## 2.7 MSRS IN THE INTEL® MICROARCHITECTURE CODE NAME NEHALEM

Table 2-14 lists model-specific registers (MSRs) that are common for Intel® microarchitecture code name Nehalem. These include Intel Core i7 and i5 processor family. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_1AH, 06\_1EH, 06\_1FH, 06\_2EH, see Table 2-1. Additional MSRs specific to 06\_1AH, 06\_1EH, 06\_1FH are listed in Table 2-15. Some MSRs listed in these tables are used by BIOS. More information about these MSR can be found at <http://biosbits.org>.

The column “Scope” represents the package/core/thread scope of individual bit field of an MSR. “Thread” means this bit field must be programmed on each logical processor independently. “Core” means the bit field must be programmed on each processor core independently, logical processors in the same core will be affected by change of this bit on the other logical processor in the same core. “Package” means the bit field must be programmed once for each physical package. Change of a bit filed with a package scope will affect all logical processors in that physical package.

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description                                 |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
| 0H               | 0   | IA32_P5_MC_ADDR            | Thread | See Section 2.22, “MSRs in Pentium Processors.” |
| 1H               | 1   | IA32_P5_MC_TYPE            | Thread | See Section 2.22, “MSRs in Pentium Processors.” |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Thread  | See Section 8.10.5, "Monitor/Mwait Address Range Determination" and Table 2-2.   |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Thread  | See Section 17.17, "Time-Stamp Counter," and see Table 2-2.  |
| 17H              | 23  | IA32_PLATFORM_ID           | Package | Platform ID (R)<br>See Table 2-2.  |
| 17H              | 23  | MSR_PLATFORM_ID            | Package | Model Specific Platform ID (R)   |
|                  |     | 49:0                       |         | Reserved   |
|                  |     | 52:50                      |         | See Table 2-2.   |
|                  |     | 63:53                      |         | Reserved   |
| 1BH              | 27  | IA32_APIC_BASE             | Thread  | See Section 10.4.4, "Local APIC Status and Location," and Table 2-2.   |
| 34H              | 52  | MSR_SMI_COUNT              | Thread  | SMI Counter (R/O)  |
|                  |     | 31:0                       |         | SMI Count (R/O)<br>Running count of SMI events since last RESET.   |
|                  |     | 63:32                      |         | Reserved   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread  | Control Features in Intel 64Processor (R/W)<br>See Table 2-2.  |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Core    | BIOS Update Trigger Register (W)<br>See Table 2-2.   |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | Thread  | BIOS Update Signature ID (RO)<br>See Table 2-2.  |
| C1H              | 193 | IA32_PMC0                  | Thread  | Performance Counter Register<br>See Table 2-2.   |
| C2H              | 194 | IA32_PMC1                  | Thread  | Performance Counter Register<br>See Table 2-2.   |
| C3H              | 195 | IA32_PMC2                  | Thread  | Performance Counter Register<br>See Table 2-2.   |
| C4H              | 196 | IA32_PMC3                  | Thread  | Performance Counter Register<br>See Table 2-2.   |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .                         |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the frequency that invariant TSC runs at. The invariant TSC frequency can be computed by multiplying this ratio by 133.33 MHz. |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 27:16                      |         | Reserved   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled.   |
|                  |     | 29                         | Package | Programmable TDC-TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDC and TDP Limits for Turbo mode are programmable. When set to 0, indicates TDC and TDP Limits for Turbo mode are not programmable.  |
|                  |     | 39:30                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 133.33MHz.  |
|                  |     | 63:48                      |         | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0 (no package C-sate support)<br>001b: C1 (Behavior is the same as 000b)<br>010b: C3<br>011b: C6<br>100b: C7<br>101b and 110b: Reserved<br>111: No package C-state limit.<br>Note: This field cannot be used to limit package C-state to C3. |
|                  |     | 9:3                        |         | Reserved   |
|                  |     | 10                         |         | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.  |
|                  |     | 14:11                      |         | Reserved   |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 15                         |        | CFG Lock (R/W)<br>When set, locks bits 15:0 of this register until next reset.   |
|                  |     | 23:16                      |        | Reserved   |
|                  |     | 24                         |        | Interrupt filtering enable (R/W)<br>When set, processor cores in a deep C-State will wake only when the event message is destined for that core. When 0, all processor cores in a deep C-State will wake for an event message.   |
|                  |     | 25                         |        | C3 state auto demotion enable (R/W)<br>When set, the processor will conditionally demote C6/C7 requests to C3 based on uncore auto-demote information.   |
|                  |     | 26                         |        | C1 state auto demotion enable (R/W)<br>When set, the processor will conditionally demote C3/C6/C7 requests to C1 based on uncore auto-demote information.  |
|                  |     | 27                         |        | Enable C3 Undemotion (R/W)   |
|                  |     | 28                         |        | Enable C1 Undemotion (R/W)   |
|                  |     | 29                         |        | Package C State Demotion Enable (R/W)  |
|                  |     | 30                         |        | Package C State UnDemotion Enable (R/W)  |
|                  |     | 63:31                      |        | Reserved   |
|                  |     | E4H                        | 228    | MSR_PMG_IO_CAPTURE_BASE  |
| 15:0             |     |                            |        | LVL_2 Base Address (R/W)<br>Specifies the base address visible to software for IO redirection. If IO MWAIT Redirection is enabled, reads to this address will be consumed by the power management logic and decoded to MWAIT instructions. When IO port address redirection is enabled, this is the IO port address reported to the OS/software. |
| 18:16            |     |                            |        | C-state Range (R/W)<br>Specifies the encoding value of the maximum C-State code name to be included when IO read to MWAIT redirection is enabled by MSR_PKG_CST_CONFIG_CONTROL[bit10]:<br>000b - C3 is the max C-State to include.<br>001b - C6 is the max C-State to include.<br>010b - C7 is the max C-State to include.                       |
| 63:19            |     |                            |        | Reserved   |
| E7H              | 231 | IA32_MPERF                 | Thread | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.   |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| E8H              | 232 | IA32_APERF                 | Thread | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP               | Thread | See Table 2-2.   |
| 174H             | 372 | IA32_SYSENTER_CS           | Thread | See Table 2-2.   |
| 175H             | 373 | IA32_SYSENTER_ESP          | Thread | See Table 2-2.   |
| 176H             | 374 | IA32_SYSENTER_EIP          | Thread | See Table 2-2.   |
| 179H             | 377 | IA32_MCG_CAP               | Thread | See Table 2-2.   |
| 17AH             | 378 | IA32_MCG_STATUS            | Thread | Global Machine Check Status  |
|                  |     | 0                          |        | RIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If cleared, the program cannot be reliably restarted.                     |
|                  |     | 1                          |        | EIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |
|                  |     | 2                          |        | MCIP<br>When set, bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception. |
|                  |     | 63:3                       |        | Reserved   |
| 186H             | 390 | IA32_PERFVTSELO            | Thread | See Table 2-2.   |
|                  |     | 7:0                        |        | Event Select   |
|                  |     | 15:8                       |        | UMask  |
|                  |     | 16                         |        | USR  |
|                  |     | 17                         |        | OS   |
|                  |     | 18                         |        | Edge   |
|                  |     | 19                         |        | PC   |
|                  |     | 20                         |        | INT  |
|                  |     | 21                         |        | AnyThread  |
|                  |     | 22                         |        | EN   |
|                  |     | 23                         |        | INV  |
|                  |     | 31:24                      |        | CMASK  |
|                  |     | 63:32                      |        | Reserved   |
| 187H             | 391 | IA32_PERFVTSEL1            | Thread | See Table 2-2.   |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |        | Register Name / Bit Fields               | Scope   | Bit Description  |
|------------------|--------|--|---------|--|
| Hex              | Dec    |  |         |  |
| 188H             | 392    | IA32_PERFEVTSEL2                         | Thread  | See Table 2-2.   |
| 189H             | 393    | IA32_PERFEVTSEL3                         | Thread  | See Table 2-2.   |
| 198H             | 408    | IA32_PERF_STATUS                         | Core    | See Table 2-2.   |
|                  |        | 15:0                                     |         | Current Performance State Value.   |
|                  |        | 63:16                                    |         | Reserved   |
| 199H             | 409    | IA32_PERF_CTL                            | Thread  | See Table 2-2.   |
| 19AH             | 410    | IA32_CLOCK_MODULATION                    | Thread  | Clock Modulation (R/W)<br>See Table 2-2.<br>IA32_CLOCK_MODULATION MSR was originally named IA32_THERM_CONTROL MSR. |
|                  |        | 0  |         | Reserved   |
|                  |        | 3:1                                      |         | On demand Clock Modulation Duty Cycle (R/W)  |
|                  |        | 4  |         | On demand Clock Modulation Enable (R/W)  |
|                  |        | 63:5                                     |         | Reserved   |
| 19BH             | 411    | IA32_THERM_INTERRUPT                     | Core    | Thermal Interrupt Control (R/W)<br>See Table 2-2.  |
| 19CH             | 412    | IA32_THERM_STATUS                        | Core    | Thermal Monitor Status (R/W)<br>See Table 2-2.   |
| 1A0H             | 416    | IA32_MISC_ENABLE                         |         | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.       |
|                  |        | 0  | Thread  | Fast-Strings Enable<br>See Table 2-2.  |
|                  |        | 2:1                                      |         | Reserved   |
|                  |        | 3  | Thread  | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2. Default value is 1.                               |
|                  |        | 6:4                                      |         | Reserved   |
|                  |        | 7  | Thread  | Performance Monitoring Available (R)<br>See Table 2-2.   |
|                  |        | 10:8                                     |         | Reserved   |
|                  |        | 11                                       | Thread  | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.  |
|                  |        | 12                                       | Thread  | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.  |
|                  |        | 15:13                                    |         | Reserved   |
|                  |        | 16                                       | Package | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.   |
| 18               | Thread | ENABLE MONITOR FSM. (R/W) See Table 2-2. |         |  |



Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 21:19                      |         | Reserved   |
|                  |     | 22                         | Thread  | Limit CPUID Maxval (R/W)<br>See Table 2-2.   |
|                  |     | 23                         | Thread  | xTPR Message Disable (R/W)<br>See Table 2-2.   |
|                  |     | 33:24                      |         | Reserved   |
|                  |     | 34                         | Thread  | XD Bit Disable (R/W)<br>See Table 2-2.   |
|                  |     | 37:35                      |         | Reserved   |
|                  |     | 38                         | Package | Turbo Mode Disable (R/W)<br>When set to 1 on processors that support Intel Turbo Boost Technology, the turbo mode feature is disabled and the IDA_Enable feature flag will be clear (CPUID.06H: EAX[1]=0).<br>When set to a 0 on processors that support IDA, CPUID.06H: EAX[1] reports the processor's support of turbo mode is enabled.<br>Note: The power-on default value is used by BIOS to detect hardware support of turbo mode. If the power-on default value is 1, turbo mode is available in the processor. If the power-on default value is 0, turbo mode is not available. |
|                  |     | 63:39                      |         | Reserved   |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Thread  | Temperature Target   |
|                  |     | 15:0                       |         | Reserved   |
|                  |     | 23:16                      |         | Temperature Target (R)<br>The minimum temperature at which PROCHOT# will be asserted. The value is degrees C.  |
|                  |     | 63:24                      |         | Reserved   |
| 1A4H             | 420 | MSR_MISC_FEATURE_CONTROL   |         | Miscellaneous Feature Control (R/W)  |
|                  |     | 0                          | Core    | L2 Hardware Prefetcher Disable (R/W)<br>If 1, disables the L2 hardware prefetcher, which fetches additional lines of code or data into the L2 cache.   |
|                  |     | 1                          | Core    | L2 Adjacent Cache Line Prefetcher Disable (R/W)<br>If 1, disables the adjacent cache line prefetcher, which fetches the cache line that comprises a cache line pair (128 bytes).   |
|                  |     | 2                          | Core    | DCU Hardware Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache prefetcher, which fetches the next cache line into L1 data cache.  |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |     | Register Name / Bit Fields    | Scope   | Bit Description  |
|------------------|-----|-------------------------------|---------|--|
| Hex              | Dec |                               |         |  |
|                  |     | 3                             | Core    | DCU IP Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache IP prefetcher, which uses sequential load history (based on instruction pointer of previous loads) to determine whether to prefetch additional lines.                      |
|                  |     | 63:4                          |         | Reserved   |
| 1A6H             | 422 | MSR_OFFCORE_RSP_0             | Thread  | Offcore Response Event Select Register (R/W)   |
| 1AAH             | 426 | MSR_MISC_PWR_MGMT             |         | Miscellaneous Power Management Control<br>Various model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 0                             | Package | EIST Hardware Coordination Disable (R/W)<br>When 0, enables hardware coordination of Enhanced Intel Speedstep Technology request from processor cores. When 1, disables hardware coordination of Enhanced Intel Speedstep Technology requests. |
|                  |     | 1                             | Thread  | Energy/Performance Bias Enable (R/W)<br>This bit makes the IA32_ENERGY_PERF_BIAS register (MSR 1B0h) visible to software with Ring 0 privileges. This bit's status (1 or 0) is also reflected by CPUID.(EAX=06h):ECX[3].                       |
|                  |     | 63:2                          |         | Reserved   |
| 1ACH             | 428 | MSR_TURBO_POWER_CURRENT_LIMIT |         | See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 14:0                          | Package | TDP Limit (R/W)<br>TDP limit in 1/8 Watt granularity.  |
|                  |     | 15                            | Package | TDP Limit Override Enable (R/W)<br>A value = 0 indicates override is not active; a value = 1 indicates override is active.   |
|                  |     | 30:16                         | Package | TDC Limit (R/W)<br>TDC limit in 1/8 Amp granularity.   |
|                  |     | 31                            | Package | TDC Limit Override Enable (R/W)<br>A value = 0 indicates override is not active; a value = 1 indicates override is active.   |
|                  |     | 63:32                         |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT         | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |     | 7:0                           | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |     | 15:8                          | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |
|                  |     | 63:32                      |         | Reserved   |
| 1C8H             | 456 | MSR_LBR_SELECT             | Core    | Last Branch Record Filtering Select Register (R/W)<br>See Section 17.9.2, "Filtering of Last Branch Records."  |
|                  |     | 0                          |         | CPL_EQ_0   |
|                  |     | 1                          |         | CPL_NEQ_0  |
|                  |     | 2                          |         | JCC  |
|                  |     | 3                          |         | NEAR_REL_CALL  |
|                  |     | 4                          |         | NEAR_IND_CALL  |
|                  |     | 5                          |         | NEAR_RET   |
|                  |     | 6                          |         | NEAR_IND_JMP   |
|                  |     | 7                          |         | NEAR_REL_JMP   |
|                  |     | 8                          |         | FAR_BRANCH   |
|                  |     | 63:9                       |         | Reserved   |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS         | Thread  | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-3) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_0_FROM_IP (at 680H).   |
| 1D9H             | 473 | IA32_DEBUGCTL              | Thread  | Debug Control (R/W)<br>See Table 2-2.  |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Thread  | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.                       |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Thread  | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled. |
| 1F2H             | 498 | IA32_SMRR_PHYSBASE         | Core    | See Table 2-2.   |
| 1F3H             | 499 | IA32_SMRR_PHYSMASK         | Core    | See Table 2-2.   |
| 1FCH             | 508 | MSR_POWER_CTL              | Core    | Power Control Register<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 0                          |         | Reserved   |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 1                          | Package | C1E Enable (R/W)<br>When set to '1', will enable the CPU to switch to the Minimum Enhanced Intel SpeedStep Technology operating point when all execution cores enter MWAIT (C1). |
|                  |     | 63:2                       |         | Reserved   |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Thread  | See Table 2-2.   |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Thread  | See Table 2-2.   |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Thread  | See Table 2-2.   |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Thread  | See Table 2-2.   |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Thread  | See Table 2-2.   |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Thread  | See Table 2-2.   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Thread  | See Table 2-2.   |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Thread  | See Table 2-2.   |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Thread  | See Table 2-2.   |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Thread  | See Table 2-2.   |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Thread  | See Table 2-2.   |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Thread  | See Table 2-2.   |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Thread  | See Table 2-2.   |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6        | Thread  | See Table 2-2.   |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7        | Thread  | See Table 2-2.   |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7        | Thread  | See Table 2-2.   |
| 210H             | 528 | IA32_MTRR_PHYSBASE8        | Thread  | See Table 2-2.   |
| 211H             | 529 | IA32_MTRR_PHYSMASK8        | Thread  | See Table 2-2.   |
| 212H             | 530 | IA32_MTRR_PHYSBASE9        | Thread  | See Table 2-2.   |
| 213H             | 531 | IA32_MTRR_PHYSMASK9        | Thread  | See Table 2-2.   |
| 250H             | 592 | IA32_MTRR_FIX64K_00000     | Thread  | See Table 2-2.   |
| 258H             | 600 | IA32_MTRR_FIX16K_80000     | Thread  | See Table 2-2.   |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000     | Thread  | See Table 2-2.   |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000      | Thread  | See Table 2-2.   |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000      | Thread  | See Table 2-2.   |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000      | Thread  | See Table 2-2.   |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000      | Thread  | See Table 2-2.   |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000      | Thread  | See Table 2-2.   |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000      | Thread  | See Table 2-2.   |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000      | Thread  | See Table 2-2.   |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000      | Thread  | See Table 2-2.   |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 277H             | 631 | IA32_PAT                   | Thread  | See Table 2-2.  |
| 280H             | 640 | IA32_MC0_CTL2              | Package | See Table 2-2.  |
| 281H             | 641 | IA32_MC1_CTL2              | Package | See Table 2-2.  |
| 282H             | 642 | IA32_MC2_CTL2              | Core    | See Table 2-2.  |
| 283H             | 643 | IA32_MC3_CTL2              | Core    | See Table 2-2.  |
| 284H             | 644 | IA32_MC4_CTL2              | Core    | See Table 2-2.  |
| 285H             | 645 | IA32_MC5_CTL2              | Core    | See Table 2-2.  |
| 286H             | 646 | IA32_MC6_CTL2              | Package | See Table 2-2.  |
| 287H             | 647 | IA32_MC7_CTL2              | Package | See Table 2-2.  |
| 288H             | 648 | IA32_MC8_CTL2              | Package | See Table 2-2.  |
| 2FFH             | 767 | IA32_MTRR_DEF_TYPE         | Thread  | Default Memory Types (R/W)<br>See Table 2-2.  |
| 309H             | 777 | IA32_FIXED_CTR0            | Thread  | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2.   |
| 30AH             | 778 | IA32_FIXED_CTR1            | Thread  | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2.   |
| 30BH             | 779 | IA32_FIXED_CTR2            | Thread  | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2.   |
| 345H             | 837 | IA32_PERF_CAPABILITIES     | Thread  | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."   |
|                  |     | 5:0                        |         | LBR Format<br>See Table 2-2.  |
|                  |     | 6                          |         | PEBS Record Format  |
|                  |     | 7                          |         | PEBSSaveArchRegs<br>See Table 2-2.  |
|                  |     | 11:8                       |         | PEBS_REC_FORMAT<br>See Table 2-2.   |
|                  |     | 12                         |         | SMM_FREEZE<br>See Table 2-2.  |
|                  |     | 63:13                      |         | Reserved  |
| 38DH             | 909 | IA32_FIXED_CTR_CTRL        | Thread  | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.   |
| 38EH             | 910 | IA32_PERF_GLOBAL_STATUS    | Thread  | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."                                     |
| 38EH             | 910 | MSR_PERF_GLOBAL_STATUS     | Thread  | Provides single-bit status used by software to query the overflow condition of each performance counter. (RO) |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 61                         |         | UNC_Ovf<br>Uncore overflowed if 1.   |
| 38FH             | 911  | IA32_PERF_GLOBAL_CTRL      | Thread  | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities."  |
| 390H             | 912  | IA32_PERF_GLOBAL_OVF_CTRL  | Thread  | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities." Allows software to clear counter overflow conditions on any combination of fixed-function PMCs (MSR_PERF_FIXED_CTRx) or general-purpose PMCs via a single WRMSR. |
| 390H             | 912  | MSR_PERF_GLOBAL_OVF_CTRL   | Thread  | (R/W)  |
|                  |      | 61                         |         | CLR_UNC_Ovf<br>Set 1 to clear UNC_Ovf.   |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Thread  | See Section 18.3.1.1.1, "Processor Event Based Sampling (PEBS)."   |
|                  |      | 0                          |         | Enable PEBS on IA32_PMC0 (R/W)   |
|                  |      | 1                          |         | Enable PEBS on IA32_PMC1 (R/W)   |
|                  |      | 2                          |         | Enable PEBS on IA32_PMC2 (R/W)   |
|                  |      | 3                          |         | Enable PEBS on IA32_PMC3 (R/W)   |
|                  |      | 31:4                       |         | Reserved   |
|                  |      | 32                         |         | Enable Load Latency on IA32_PMC0 (R/W)   |
|                  |      | 33                         |         | Enable Load Latency on IA32_PMC1 (R/W)   |
|                  |      | 34                         |         | Enable Load Latency on IA32_PMC2 (R/W)   |
|                  |      | 35                         |         | Enable Load Latency on IA32_PMC3 (R/W)   |
|                  |      | 63:36                      |         | Reserved   |
| 3F6H             | 1014 | MSR_PEBS_LD_LAT            | Thread  | See Section 18.3.1.1.2, "Load Latency Performance Monitoring Facility."  |
|                  |      | 15:0                       |         | Minimum threshold latency value of tagged load operation that will be counted. (R/W)   |
|                  |      | 63:36                      |         | Reserved   |
| 3F8H             | 1016 | MSR_PKG_C3_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |
|                  |      | 63:0                       |         | Package C3 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C3 states. Count at the same frequency as the TSC.   |
| 3F9H             | 1017 | MSR_PKG_C6_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:0                       |         | Package C6 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C6 states. Count at the same frequency as the TSC.  |
| 3FAH             | 1018 | MSR_PKG_C7_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.   |
|                  |      | 63:0                       |         | Package C7 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C7 states. Count at the same frequency as the TSC.  |
| 3FCH             | 1020 | MSR_CORE_C3_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.   |
|                  |      | 63:0                       |         | CORE C3 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C3 states. Count at the same frequency as the TSC.  |
| 3FDH             | 1021 | MSR_CORE_C6_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.   |
|                  |      | 63:0                       |         | CORE C6 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C6 states. Count at the same frequency as the TSC.  |
| 400H             | 1024 | IA32_MCO_CTL               | Package | See Section 15.3.2.1, "IA32_MCI_CTL MSRs."  |
| 401H             | 1025 | IA32_MCO_STATUS            | Package | See Section 15.3.2.2, "IA32_MCI_STATUS MSRs."   |
| 402H             | 1026 | IA32_MCO_ADDR              | Package | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 403H             | 1027 | IA32_MCO_MISC              | Package | See Section 15.3.2.4, "IA32_MCI_MISC MSRs."   |
| 404H             | 1028 | IA32_MC1_CTL               | Package | See Section 15.3.2.1, "IA32_MCI_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS            | Package | See Section 15.3.2.2, "IA32_MCI_STATUS MSRs."   |
| 406H             | 1030 | IA32_MC1_ADDR              | Package | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs."<br>The IA32_MC1_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC1_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |

**Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 407H             | 1031 | IA32_MC1_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 408H             | 1032 | IA32_MC2_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 409H             | 1033 | IA32_MC2_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40BH             | 1035 | IA32_MC2_MISC              | Core    | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 40CH             | 1036 | IA32_MC3_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 40DH             | 1037 | IA32_MC3_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40EH             | 1038 | IA32_MC3_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 40FH             | 1039 | IA32_MC3_MISC              | Core    | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 410H             | 1040 | IA32_MC4_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC4_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC4_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 413H             | 1043 | IA32_MC4_MISC              | Core    | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 414H             | 1044 | IA32_MC5_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 415H             | 1045 | IA32_MC5_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 416H             | 1046 | IA32_MC5_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 417H             | 1047 | IA32_MC5_MISC              | Core    | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 419H             | 1049 | IA32_MC6_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.  |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 41BH             | 1051 | IA32_MC6_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."   |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |



Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package | See Section 15.3.2.2, "IA32_MCI_STATUS MSRS" and Chapter 16.   |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs."  |
| 41FH             | 1055 | IA32_MC7_MISC              | Package | See Section 15.3.2.4, "IA32_MCI_MISC MSRs."  |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCI_CTL MSRs."   |
| 421H             | 1057 | IA32_MC8_STATUS            | Package | See Section 15.3.2.2, "IA32_MCI_STATUS MSRS" and Chapter 16.   |
| 422H             | 1058 | IA32_MC8_ADDR              | Package | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs."  |
| 423H             | 1059 | IA32_MC8_MISC              | Package | See Section 15.3.2.4, "IA32_MCI_MISC MSRs."  |
| 480H             | 1152 | IA32_VMX_BASIC             | Thread  | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."                     |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL      | Thread  | Capability Reporting Register of Pin-based VM-execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls." |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL     | Thread  | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."     |
| 483H             | 1155 | IA32_VMX_EXIT_CTL          | Thread  | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Table 2-2.<br>See Appendix A.4, "VM-Exit Controls."                     |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL         | Thread  | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Table 2-2.<br>See Appendix A.5, "VM-Entry Controls."                   |
| 485H             | 1157 | IA32_VMX_MISC              | Thread  | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.6, "Miscellaneous Data."                |
| 486H             | 1158 | IA32_VMX_CRO_FIXED0        | Thread  | Capability Reporting Register of CRO Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CRO."             |
| 487H             | 1159 | IA32_VMX_CRO_FIXED1        | Thread  | Capability Reporting Register of CRO Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CRO."             |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0        | Thread | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."  |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1        | Thread | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."  |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM         | Thread | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.<br>See Appendix A.9, "VMCS Enumeration."  |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL2    | Thread | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."  |
| 600H             | 1536 | IA32_DS_AREA               | Thread | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."   |
| 680H             | 1664 | MSR_LASTBRANCH_0_FROM_IP   | Thread | Last Branch Record 0 From IP (R/W)<br>One of sixteen pairs of last branch record registers on the last branch record stack. The From_IP part of the stack contains pointers to the source instruction. See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.9.1 and record format in Section 17.4.8.1.</li> </ul> |
| 681H             | 1665 | MSR_LASTBRANCH_1_FROM_IP   | Thread | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 682H             | 1666 | MSR_LASTBRANCH_2_FROM_IP   | Thread | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 683H             | 1667 | MSR_LASTBRANCH_3_FROM_IP   | Thread | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 684H             | 1668 | MSR_LASTBRANCH_4_FROM_IP   | Thread | Last Branch Record 4 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 685H             | 1669 | MSR_LASTBRANCH_5_FROM_IP   | Thread | Last Branch Record 5 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 686H             | 1670 | MSR_LASTBRANCH_6_FROM_IP   | Thread | Last Branch Record 6 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 687H             | 1671 | MSR_LASTBRANCH_7_FROM_IP   | Thread | Last Branch Record 7 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 688H             | 1672 | MSR_LASTBRANCH_8_FROM_IP   | Thread | Last Branch Record 8 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| 689H             | 1673 | MSR_LASTBRANCH_9_FROM_IP   | Thread | Last Branch Record 9 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 68AH             | 1674 | MSR_LASTBRANCH_10_FROM_IP  | Thread | Last Branch Record 10 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68BH             | 1675 | MSR_LASTBRANCH_11_FROM_IP  | Thread | Last Branch Record 11 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68CH             | 1676 | MSR_LASTBRANCH_12_FROM_IP  | Thread | Last Branch Record 12 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68DH             | 1677 | MSR_LASTBRANCH_13_FROM_IP  | Thread | Last Branch Record 13 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68EH             | 1678 | MSR_LASTBRANCH_14_FROM_IP  | Thread | Last Branch Record 14 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68FH             | 1679 | MSR_LASTBRANCH_15_FROM_IP  | Thread | Last Branch Record 15 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 6C0H             | 1728 | MSR_LASTBRANCH_0_TO_IP     | Thread | Last Branch Record 0 To IP (R/W)<br>One of sixteen pairs of last branch record registers on the last branch record stack. This part of the stack contains pointers to the destination instruction. |
| 6C1H             | 1729 | MSR_LASTBRANCH_1_TO_IP     | Thread | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C2H             | 1730 | MSR_LASTBRANCH_2_TO_IP     | Thread | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C3H             | 1731 | MSR_LASTBRANCH_3_TO_IP     | Thread | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C4H             | 1732 | MSR_LASTBRANCH_4_TO_IP     | Thread | Last Branch Record 4 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C5H             | 1733 | MSR_LASTBRANCH_5_TO_IP     | Thread | Last Branch Record 5 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C6H             | 1734 | MSR_LASTBRANCH_6_TO_IP     | Thread | Last Branch Record 6 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C7H             | 1735 | MSR_LASTBRANCH_7_TO_IP     | Thread | Last Branch Record 7 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C8H             | 1736 | MSR_LASTBRANCH_8_TO_IP     | Thread | Last Branch Record 8 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C9H             | 1737 | MSR_LASTBRANCH_9_TO_IP     | Thread | Last Branch Record 9 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6CAH             | 1738 | MSR_LASTBRANCH_10_TO_IP    | Thread | Last Branch Record 10 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 6CBH             | 1739 | MSR_LASTBRANCH_11_TO_IP    | Thread | Last Branch Record 11 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CCH             | 1740 | MSR_LASTBRANCH_12_TO_IP    | Thread | Last Branch Record 12 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CDH             | 1741 | MSR_LASTBRANCH_13_TO_IP    | Thread | Last Branch Record 13 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CEH             | 1742 | MSR_LASTBRANCH_14_TO_IP    | Thread | Last Branch Record 14 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 6CFH             | 1743 | MSR_LASTBRANCH_15_TO_IP    | Thread | Last Branch Record 15 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP. |
| 802H             | 2050 | IA32_X2APIC_APICID         | Thread | x2APIC ID Register (R/O)  |
| 803H             | 2051 | IA32_X2APIC_VERSION        | Thread | x2APIC Version Register (R/O)   |
| 808H             | 2056 | IA32_X2APIC_TPR            | Thread | x2APIC Task Priority Register (R/W)   |
| 80AH             | 2058 | IA32_X2APIC_PPR            | Thread | x2APIC Processor Priority Register (R/O)  |
| 80BH             | 2059 | IA32_X2APIC_EOI            | Thread | x2APIC EOI Register (W/O)   |
| 80DH             | 2061 | IA32_X2APIC_LDR            | Thread | x2APIC Logical Destination Register (R/O)                                       |
| 80FH             | 2063 | IA32_X2APIC_SIVR           | Thread | x2APIC Spurious Interrupt Vector Register (R/W)                                 |
| 810H             | 2064 | IA32_X2APIC_ISR0           | Thread | x2APIC In-Service Register Bits [31:0] (R/O)                                    |
| 811H             | 2065 | IA32_X2APIC_ISR1           | Thread | x2APIC In-Service Register Bits [63:32] (R/O)                                   |
| 812H             | 2066 | IA32_X2APIC_ISR2           | Thread | x2APIC In-Service Register Bits [95:64] (R/O)                                   |
| 813H             | 2067 | IA32_X2APIC_ISR3           | Thread | x2APIC In-Service Register Bits [127:96] (R/O)                                  |
| 814H             | 2068 | IA32_X2APIC_ISR4           | Thread | x2APIC In-Service Register Bits [159:128] (R/O)                                 |
| 815H             | 2069 | IA32_X2APIC_ISR5           | Thread | x2APIC In-Service Register Bits [191:160] (R/O)                                 |
| 816H             | 2070 | IA32_X2APIC_ISR6           | Thread | x2APIC In-Service Register Bits [223:192] (R/O)                                 |
| 817H             | 2071 | IA32_X2APIC_ISR7           | Thread | x2APIC In-Service Register Bits [255:224] (R/O)                                 |
| 818H             | 2072 | IA32_X2APIC_TMR0           | Thread | x2APIC Trigger Mode Register Bits [31:0] (R/O)                                  |
| 819H             | 2073 | IA32_X2APIC_TMR1           | Thread | x2APIC Trigger Mode Register Bits [63:32] (R/O)                                 |
| 81AH             | 2074 | IA32_X2APIC_TMR2           | Thread | x2APIC Trigger Mode Register Bits [95:64] (R/O)                                 |
| 81BH             | 2075 | IA32_X2APIC_TMR3           | Thread | x2APIC Trigger Mode Register Bits [127:96] (R/O)                                |
| 81CH             | 2076 | IA32_X2APIC_TMR4           | Thread | x2APIC Trigger Mode Register Bits [159:128] (R/O)                               |
| 81DH             | 2077 | IA32_X2APIC_TMR5           | Thread | x2APIC Trigger Mode Register Bits [191:160] (R/O)                               |
| 81EH             | 2078 | IA32_X2APIC_TMR6           | Thread | x2APIC Trigger Mode Register Bits [223:192] (R/O)                               |
| 81FH             | 2079 | IA32_X2APIC_TMR7           | Thread | x2APIC Trigger Mode Register Bits [255:224] (R/O)                               |
| 820H             | 2080 | IA32_X2APIC_IRR0           | Thread | x2APIC Interrupt Request Register Bits [31:0] (R/O)                             |
| 821H             | 2081 | IA32_X2APIC_IRR1           | Thread | x2APIC Interrupt Request Register Bits [63:32] (R/O)                            |
| 822H             | 2082 | IA32_X2APIC_IRR2           | Thread | x2APIC Interrupt Request Register Bits [95:64] (R/O)                            |

Table 2-14. MSRs in Processors Based on Intel® Microarchitecture Code Name Nehalem (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 823H             | 2083 | IA32_X2APIC_IRR3           | Thread | x2APIC Interrupt Request Register Bits [127:96] (R/O)   |
| 824H             | 2084 | IA32_X2APIC_IRR4           | Thread | x2APIC Interrupt Request Register Bits [159:128] (R/O)  |
| 825H             | 2085 | IA32_X2APIC_IRR5           | Thread | x2APIC Interrupt Request Register Bits [191:160] (R/O)  |
| 826H             | 2086 | IA32_X2APIC_IRR6           | Thread | x2APIC Interrupt Request Register Bits [223:192] (R/O)  |
| 827H             | 2087 | IA32_X2APIC_IRR7           | Thread | x2APIC Interrupt Request Register Bits [255:224] (R/O)  |
| 828H             | 2088 | IA32_X2APIC_ESR            | Thread | x2APIC Error Status Register (R/W)  |
| 82FH             | 2095 | IA32_X2APIC_LVT_CMCI       | Thread | x2APIC LVT Corrected Machine Check Interrupt Register (R/W)   |
| 830H             | 2096 | IA32_X2APIC_ICR            | Thread | x2APIC Interrupt Command Register (R/W)   |
| 832H             | 2098 | IA32_X2APIC_LVT_TIMER      | Thread | x2APIC LVT Timer Interrupt Register (R/W)   |
| 833H             | 2099 | IA32_X2APIC_LVT_THERMAL    | Thread | x2APIC LVT Thermal Sensor Interrupt Register (R/W)  |
| 834H             | 2100 | IA32_X2APIC_LVT_PMI        | Thread | x2APIC LVT Performance Monitor Register (R/W)   |
| 835H             | 2101 | IA32_X2APIC_LVT_LINT0      | Thread | x2APIC LVT LINT0 Register (R/W)   |
| 836H             | 2102 | IA32_X2APIC_LVT_LINT1      | Thread | x2APIC LVT LINT1 Register (R/W)   |
| 837H             | 2103 | IA32_X2APIC_LVT_ERROR      | Thread | x2APIC LVT Error Register (R/W)   |
| 838H             | 2104 | IA32_X2APIC_INIT_COUNT     | Thread | x2APIC Initial Count Register (R/W)   |
| 839H             | 2105 | IA32_X2APIC_CUR_COUNT      | Thread | x2APIC Current Count Register (R/O)   |
| 83EH             | 2110 | IA32_X2APIC_DIV_CONF       | Thread | x2APIC Divide Configuration Register (R/W)  |
| 83FH             | 2111 | IA32_X2APIC_SELF_IPI       | Thread | x2APIC Self IPI Register (W/O)  |
| C000_0080H       |      | IA32_EFER                  | Thread | Extended Feature Enables<br>See Table 2-2.  |
| C000_0081H       |      | IA32_STAR                  | Thread | System Call Target Address (R/W)<br>See Table 2-2.  |
| C000_0082H       |      | IA32_LSTAR                 | Thread | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.  |
| C000_0084H       |      | IA32_FMASK                 | Thread | System Call Flag Mask (R/W)<br>See Table 2-2.   |
| C000_0100H       |      | IA32_FS_BASE               | Thread | Map of BASE Address of FS (R/W)<br>See Table 2-2.   |
| C000_0101H       |      | IA32_GS_BASE               | Thread | Map of BASE Address of GS (R/W)<br>See Table 2-2.   |
| C000_0102H       |      | IA32_KERNEL_GS_BASE        | Thread | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.   |
| C000_0103H       |      | IA32_TSC_AUX               | Thread | AUXILIARY TSC Signature (R/W)<br>See Table 2-2 and Section 17.17.2, "IA32_TSC_AUX Register and RDTSCP Support." |

### 2.7.1 Additional MSRs in the Intel® Xeon® Processor 5500 and 3400 Series

Intel Xeon Processor 5500 and 3400 series support additional model-specific registers listed in Table 2-15. These MSRs also apply to Intel Core i7 and i5 processor family CPUID signature with DisplayFamily\_DisplayModel of 06\_1AH, 06\_1EH and 06\_1FH, see Table 2-1.

**Table 2-15. Additional MSRs in Intel® Xeon® Processor 5500 and 3400 Series**

| Register Address |          | Register Name / Bit Fields      | Scope   | Bit Description  |
|------------------|----------|---------------------------------|---------|--|
| Hex              | Dec      |                                 |         |  |
| 1ADH             | 429      | MSR_TURBO_RATIO_LIMIT           | Package | Actual maximum turbo frequency is multiplied by 133.33MHz.<br>(Not available in model 06_2EH.) |
|                  |          | 7:0                             |         | Maximum Turbo Ratio Limit 1C (R/O)<br>Maximum Turbo mode ratio limit with 1 core active.       |
|                  |          | 15:8                            |         | Maximum Turbo Ratio Limit 2C (R/O)<br>Maximum Turbo mode ratio limit with 2 cores active.      |
|                  |          | 23:16                           |         | Maximum Turbo Ratio Limit 3C (R/O)<br>Maximum Turbo mode ratio limit with 3 cores active.      |
|                  |          | 31:24                           |         | Maximum Turbo Ratio Limit 4C (R/O)<br>Maximum Turbo mode ratio limit with 4 cores active.      |
|                  |          | 63:32                           |         | Reserved   |
| 301H             | 769      | MSR_GQ_SNOOP_MESF               | Package |  |
|                  |          | 0                               |         | From M to S (R/W)  |
|                  |          | 1                               |         | From E to S (R/W)  |
|                  |          | 2                               |         | From S to S (R/W)  |
|                  |          | 3                               |         | From F to S (R/W)  |
|                  |          | 4                               |         | From M to I (R/W)  |
|                  |          | 5                               |         | From E to I (R/W)  |
|                  |          | 6                               |         | From S to I (R/W)  |
|                  |          | 7                               |         | From F to I (R/W)  |
| 63:8             | Reserved |                                 |         |  |
| 391H             | 913      | MSR_UNCORE_PERF_GLOBAL_CTRL     | Package | See Section 18.3.1.2.1, "Uncore Performance Monitoring Management Facility."                   |
| 392H             | 914      | MSR_UNCORE_PERF_GLOBAL_STATUS   | Package | See Section 18.3.1.2.1, "Uncore Performance Monitoring Management Facility."                   |
| 393H             | 915      | MSR_UNCORE_PERF_GLOBAL_OVF_CTRL | Package | See Section 18.3.1.2.1, "Uncore Performance Monitoring Management Facility."                   |
| 394H             | 916      | MSR_UNCORE_FIXED_CTR0           | Package | See Section 18.3.1.2.1, "Uncore Performance Monitoring Management Facility."                   |
| 395H             | 917      | MSR_UNCORE_FIXED_CTR_CTRL       | Package | See Section 18.3.1.2.1, "Uncore Performance Monitoring Management Facility."                   |
| 396H             | 918      | MSR_UNCORE_ADDR_OPCODE_MATCH    | Package | See Section 18.3.1.2.3, "Uncore Address/Opcode Match MSR."                                     |
| 3B0H             | 960      | MSR_UNCORE_PMC0                 | Package | See Section 18.3.1.2.2, "Uncore Performance Event Configuration Facility."                     |

**Table 2-15. Additional MSRs in Intel® Xeon® Processor 5500 and 3400 Series (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 3B1H             | 961 | MSR_UNCORE_PMC1            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B2H             | 962 | MSR_UNCORE_PMC2            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B3H             | 963 | MSR_UNCORE_PMC3            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B4H             | 964 | MSR_UNCORE_PMC4            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B5H             | 965 | MSR_UNCORE_PMC5            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B6H             | 966 | MSR_UNCORE_PMC6            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3B7H             | 967 | MSR_UNCORE_PMC7            | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C0H             | 944 | MSR_UNCORE_PERFEVTSELO     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C1H             | 945 | MSR_UNCORE_PERFEVTSEL1     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C2H             | 946 | MSR_UNCORE_PERFEVTSEL2     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C3H             | 947 | MSR_UNCORE_PERFEVTSEL3     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C4H             | 948 | MSR_UNCORE_PERFEVTSEL4     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C5H             | 949 | MSR_UNCORE_PERFEVTSEL5     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C6H             | 950 | MSR_UNCORE_PERFEVTSEL6     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |
| 3C7H             | 951 | MSR_UNCORE_PERFEVTSEL7     | Package | See Section 18.3.1.2.2, “Uncore Performance Event Configuration Facility.” |

## 2.7.2 Additional MSRs in the Intel® Xeon® Processor 7500 Series

Intel Xeon Processor 7500 series support MSRs listed in Table 2-14 (except MSR address 1ADH) and additional model-specific registers listed in Table 2-16. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2EH.

**Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                                   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Reserved<br>Attempt to read/write will cause #UD. |

**Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 289H             | 649  | IA32_MC9_CTL2              | Package | See Table 2-2.   |
| 28AH             | 650  | IA32_MC10_CTL2             | Package | See Table 2-2.   |
| 28BH             | 651  | IA32_MC11_CTL2             | Package | See Table 2-2.   |
| 28CH             | 652  | IA32_MC12_CTL2             | Package | See Table 2-2.   |
| 28DH             | 653  | IA32_MC13_CTL2             | Package | See Table 2-2.   |
| 28EH             | 654  | IA32_MC14_CTL2             | Package | See Table 2-2.   |
| 28FH             | 655  | IA32_MC15_CTL2             | Package | See Table 2-2.   |
| 290H             | 656  | IA32_MC16_CTL2             | Package | See Table 2-2.   |
| 291H             | 657  | IA32_MC17_CTL2             | Package | See Table 2-2.   |
| 292H             | 658  | IA32_MC18_CTL2             | Package | See Table 2-2.   |
| 293H             | 659  | IA32_MC19_CTL2             | Package | See Table 2-2.   |
| 294H             | 660  | IA32_MC20_CTL2             | Package | See Table 2-2.   |
| 295H             | 661  | IA32_MC21_CTL2             | Package | See Table 2-2.   |
| 394H             | 816  | MSR_W_PMON_FIXED_CTR       | Package | Uncore W-box perfmon fixed counter.                          |
| 395H             | 817  | MSR_W_PMON_FIXED_CTR_CTL   | Package | Uncore U-box perfmon fixed counter control MSR.              |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 425H             | 1061 | IA32_MC9_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 426H             | 1062 | IA32_MC9_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 427H             | 1063 | IA32_MC9_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 429H             | 1065 | IA32_MC10_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 42BH             | 1067 | IA32_MC10_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 42FH             | 1071 | IA32_MC11_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 431H             | 1073 | IA32_MC12_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 432H             | 1074 | IA32_MC12_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 433H             | 1075 | IA32_MC12_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 435H             | 1077 | IA32_MC13_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |



Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 436H             | 1078 | IA32_MC13_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 437H             | 1079 | IA32_MC13_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 439H             | 1081 | IA32_MC14_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 43BH             | 1083 | IA32_MC14_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 43FH             | 1087 | IA32_MC15_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 441H             | 1089 | IA32_MC16_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 442H             | 1090 | IA32_MC16_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 443H             | 1091 | IA32_MC16_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 445H             | 1093 | IA32_MC17_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 446H             | 1094 | IA32_MC17_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 447H             | 1095 | IA32_MC17_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 449H             | 1097 | IA32_MC18_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 44BH             | 1099 | IA32_MC18_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 44FH             | 1103 | IA32_MC19_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 450H             | 1104 | IA32_MC20_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 451H             | 1105 | IA32_MC20_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 452H             | 1106 | IA32_MC20_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 453H             | 1107 | IA32_MC20_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 454H             | 1108 | IA32_MC21_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 455H             | 1109 | IA32_MC21_STATUS           | Package | See Section 15.3.2.2, "IA32_MCI_STATUS MSRS" and Chapter 16. |
| 456H             | 1110 | IA32_MC21_ADDR             | Package | See Section 15.3.2.3, "IA32_MCI_ADDR MSRS."                  |
| 457H             | 1111 | IA32_MC21_MISC             | Package | See Section 15.3.2.4, "IA32_MCI_MISC MSRS."                  |
| C00H             | 3072 | MSR_U_PMON_GLOBAL_CTRL     | Package | Uncore U-box perfmon global control MSR.                     |
| C01H             | 3073 | MSR_U_PMON_GLOBAL_STATUS   | Package | Uncore U-box perfmon global status MSR.                      |
| C02H             | 3074 | MSR_U_PMON_GLOBAL_OVF_CTRL | Package | Uncore U-box perfmon global overflow control MSR.            |
| C10H             | 3088 | MSR_U_PMON_EVNT_SEL        | Package | Uncore U-box perfmon event select MSR.                       |
| C11H             | 3089 | MSR_U_PMON_CTR             | Package | Uncore U-box perfmon counter MSR.                            |
| C20H             | 3104 | MSR_B0_PMON_BOX_CTRL       | Package | Uncore B-box 0 perfmon local box control MSR.                |
| C21H             | 3105 | MSR_B0_PMON_BOX_STATUS     | Package | Uncore B-box 0 perfmon local box status MSR.                 |
| C22H             | 3106 | MSR_B0_PMON_BOX_OVF_CTRL   | Package | Uncore B-box 0 perfmon local box overflow control MSR.       |
| C30H             | 3120 | MSR_B0_PMON_EVNT_SELO      | Package | Uncore B-box 0 perfmon event select MSR.                     |
| C31H             | 3121 | MSR_B0_PMON_CTR0           | Package | Uncore B-box 0 perfmon counter MSR.                          |
| C32H             | 3122 | MSR_B0_PMON_EVNT_SEL1      | Package | Uncore B-box 0 perfmon event select MSR.                     |
| C33H             | 3123 | MSR_B0_PMON_CTR1           | Package | Uncore B-box 0 perfmon counter MSR.                          |
| C34H             | 3124 | MSR_B0_PMON_EVNT_SEL2      | Package | Uncore B-box 0 perfmon event select MSR.                     |
| C35H             | 3125 | MSR_B0_PMON_CTR2           | Package | Uncore B-box 0 perfmon counter MSR.                          |
| C36H             | 3126 | MSR_B0_PMON_EVNT_SEL3      | Package | Uncore B-box 0 perfmon event select MSR.                     |
| C37H             | 3127 | MSR_B0_PMON_CTR3           | Package | Uncore B-box 0 perfmon counter MSR.                          |
| C40H             | 3136 | MSR_S0_PMON_BOX_CTRL       | Package | Uncore S-box 0 perfmon local box control MSR.                |
| C41H             | 3137 | MSR_S0_PMON_BOX_STATUS     | Package | Uncore S-box 0 perfmon local box status MSR.                 |
| C42H             | 3138 | MSR_S0_PMON_BOX_OVF_CTRL   | Package | Uncore S-box 0 perfmon local box overflow control MSR.       |
| C50H             | 3152 | MSR_S0_PMON_EVNT_SELO      | Package | Uncore S-box 0 perfmon event select MSR.                     |
| C51H             | 3153 | MSR_S0_PMON_CTR0           | Package | Uncore S-box 0 perfmon counter MSR.                          |
| C52H             | 3154 | MSR_S0_PMON_EVNT_SEL1      | Package | Uncore S-box 0 perfmon event select MSR.                     |
| C53H             | 3155 | MSR_S0_PMON_CTR1           | Package | Uncore S-box 0 perfmon counter MSR.                          |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| C54H             | 3156 | MSR_S0_PMON_EVNT_SEL2      | Package | Uncore S-box 0 perfmon event select MSR.               |
| C55H             | 3157 | MSR_S0_PMON_CTR2           | Package | Uncore S-box 0 perfmon counter MSR.                    |
| C56H             | 3158 | MSR_S0_PMON_EVNT_SEL3      | Package | Uncore S-box 0 perfmon event select MSR.               |
| C57H             | 3159 | MSR_S0_PMON_CTR3           | Package | Uncore S-box 0 perfmon counter MSR.                    |
| C60H             | 3168 | MSR_B1_PMON_BOX_CTRL       | Package | Uncore B-box 1 perfmon local box control MSR.          |
| C61H             | 3169 | MSR_B1_PMON_BOX_STATUS     | Package | Uncore B-box 1 perfmon local box status MSR.           |
| C62H             | 3170 | MSR_B1_PMON_BOX_OVF_CTRL   | Package | Uncore B-box 1 perfmon local box overflow control MSR. |
| C70H             | 3184 | MSR_B1_PMON_EVNT_SELO      | Package | Uncore B-box 1 perfmon event select MSR.               |
| C71H             | 3185 | MSR_B1_PMON_CTR0           | Package | Uncore B-box 1 perfmon counter MSR.                    |
| C72H             | 3186 | MSR_B1_PMON_EVNT_SEL1      | Package | Uncore B-box 1 perfmon event select MSR.               |
| C73H             | 3187 | MSR_B1_PMON_CTR1           | Package | Uncore B-box 1 perfmon counter MSR.                    |
| C74H             | 3188 | MSR_B1_PMON_EVNT_SEL2      | Package | Uncore B-box 1 perfmon event select MSR.               |
| C75H             | 3189 | MSR_B1_PMON_CTR2           | Package | Uncore B-box 1 perfmon counter MSR.                    |
| C76H             | 3190 | MSR_B1_PMON_EVNT_SEL3      | Package | Uncore B-box 1 vperfmon event select MSR.              |
| C77H             | 3191 | MSR_B1_PMON_CTR3           | Package | Uncore B-box 1 perfmon counter MSR.                    |
| C80H             | 3120 | MSR_W_PMON_BOX_CTRL        | Package | Uncore W-box perfmon local box control MSR.            |
| C81H             | 3121 | MSR_W_PMON_BOX_STATUS      | Package | Uncore W-box perfmon local box status MSR.             |
| C82H             | 3122 | MSR_W_PMON_BOX_OVF_CTRL    | Package | Uncore W-box perfmon local box overflow control MSR.   |
| C90H             | 3136 | MSR_W_PMON_EVNT_SELO       | Package | Uncore W-box perfmon event select MSR.                 |
| C91H             | 3137 | MSR_W_PMON_CTR0            | Package | Uncore W-box perfmon counter MSR.                      |
| C92H             | 3138 | MSR_W_PMON_EVNT_SEL1       | Package | Uncore W-box perfmon event select MSR.                 |
| C93H             | 3139 | MSR_W_PMON_CTR1            | Package | Uncore W-box perfmon counter MSR.                      |
| C94H             | 3140 | MSR_W_PMON_EVNT_SEL2       | Package | Uncore W-box perfmon event select MSR.                 |
| C95H             | 3141 | MSR_W_PMON_CTR2            | Package | Uncore W-box perfmon counter MSR.                      |
| C96H             | 3142 | MSR_W_PMON_EVNT_SEL3       | Package | Uncore W-box perfmon event select MSR.                 |
| C97H             | 3143 | MSR_W_PMON_CTR3            | Package | Uncore W-box perfmon counter MSR.                      |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| CA0H             | 3232 | MSR_M0_PMON_BOX_CTRL       | Package | Uncore M-box 0 perfmon local box control MSR.          |
| CA1H             | 3233 | MSR_M0_PMON_BOX_STATUS     | Package | Uncore M-box 0 perfmon local box status MSR.           |
| CA2H             | 3234 | MSR_M0_PMON_BOX_OVF_CTRL   | Package | Uncore M-box 0 perfmon local box overflow control MSR. |
| CA4H             | 3236 | MSR_M0_PMON_TIMESTAMP      | Package | Uncore M-box 0 perfmon time stamp unit select MSR.     |
| CA5H             | 3237 | MSR_M0_PMON_DSP            | Package | Uncore M-box 0 perfmon DSP unit select MSR.            |
| CA6H             | 3238 | MSR_M0_PMON_ISS            | Package | Uncore M-box 0 perfmon ISS unit select MSR.            |
| CA7H             | 3239 | MSR_M0_PMON_MAP            | Package | Uncore M-box 0 perfmon MAP unit select MSR.            |
| CA8H             | 3240 | MSR_M0_PMON_MSC_THR        | Package | Uncore M-box 0 perfmon MIC THR select MSR.             |
| CA9H             | 3241 | MSR_M0_PMON_PGT            | Package | Uncore M-box 0 perfmon PGT unit select MSR.            |
| CAAH             | 3242 | MSR_M0_PMON_PLD            | Package | Uncore M-box 0 perfmon PLD unit select MSR.            |
| CABH             | 3243 | MSR_M0_PMON_ZDP            | Package | Uncore M-box 0 perfmon ZDP unit select MSR.            |
| CBOH             | 3248 | MSR_M0_PMON_EVNT_SELO      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CB1H             | 3249 | MSR_M0_PMON_CTR0           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CB2H             | 3250 | MSR_M0_PMON_EVNT_SEL1      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CB3H             | 3251 | MSR_M0_PMON_CTR1           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CB4H             | 3252 | MSR_M0_PMON_EVNT_SEL2      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CB5H             | 3253 | MSR_M0_PMON_CTR2           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CB6H             | 3254 | MSR_M0_PMON_EVNT_SEL3      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CB7H             | 3255 | MSR_M0_PMON_CTR3           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CB8H             | 3256 | MSR_M0_PMON_EVNT_SEL4      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CB9H             | 3257 | MSR_M0_PMON_CTR4           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CBAH             | 3258 | MSR_M0_PMON_EVNT_SEL5      | Package | Uncore M-box 0 perfmon event select MSR.               |
| CBBH             | 3259 | MSR_M0_PMON_CTR5           | Package | Uncore M-box 0 perfmon counter MSR.                    |
| CC0H             | 3264 | MSR_S1_PMON_BOX_CTRL       | Package | Uncore S-box 1 perfmon local box control MSR.          |
| CC1H             | 3265 | MSR_S1_PMON_BOX_STATUS     | Package | Uncore S-box 1 perfmon local box status MSR.           |
| CC2H             | 3266 | MSR_S1_PMON_BOX_OVF_CTRL   | Package | Uncore S-box 1 perfmon local box overflow control MSR. |
| CDOH             | 3280 | MSR_S1_PMON_EVNT_SELO      | Package | Uncore S-box 1 perfmon event select MSR.               |
| CD1H             | 3281 | MSR_S1_PMON_CTR0           | Package | Uncore S-box 1 perfmon counter MSR.                    |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| CD2H             | 3282 | MSR_S1_PMON_EVNT_SEL1      | Package | Uncore S-box 1 perfmon event select MSR.               |
| CD3H             | 3283 | MSR_S1_PMON_CTR1           | Package | Uncore S-box 1 perfmon counter MSR.                    |
| CD4H             | 3284 | MSR_S1_PMON_EVNT_SEL2      | Package | Uncore S-box 1 perfmon event select MSR.               |
| CD5H             | 3285 | MSR_S1_PMON_CTR2           | Package | Uncore S-box 1 perfmon counter MSR.                    |
| CD6H             | 3286 | MSR_S1_PMON_EVNT_SEL3      | Package | Uncore S-box 1 perfmon event select MSR.               |
| CD7H             | 3287 | MSR_S1_PMON_CTR3           | Package | Uncore S-box 1 perfmon counter MSR.                    |
| CE0H             | 3296 | MSR_M1_PMON_BOX_CTRL       | Package | Uncore M-box 1 perfmon local box control MSR.          |
| CE1H             | 3297 | MSR_M1_PMON_BOX_STATUS     | Package | Uncore M-box 1 perfmon local box status MSR.           |
| CE2H             | 3298 | MSR_M1_PMON_BOX_OVF_CTRL   | Package | Uncore M-box 1 perfmon local box overflow control MSR. |
| CE4H             | 3300 | MSR_M1_PMON_TIMESTAMP      | Package | Uncore M-box 1 perfmon time stamp unit select MSR.     |
| CE5H             | 3301 | MSR_M1_PMON_DSP            | Package | Uncore M-box 1 perfmon DSP unit select MSR.            |
| CE6H             | 3302 | MSR_M1_PMON_ISS            | Package | Uncore M-box 1 perfmon ISS unit select MSR.            |
| CE7H             | 3303 | MSR_M1_PMON_MAP            | Package | Uncore M-box 1 perfmon MAP unit select MSR.            |
| CE8H             | 3304 | MSR_M1_PMON_MSC_THR        | Package | Uncore M-box 1 perfmon MIC THR select MSR.             |
| CE9H             | 3305 | MSR_M1_PMON_PGT            | Package | Uncore M-box 1 perfmon PGT unit select MSR.            |
| CEAH             | 3306 | MSR_M1_PMON_PLD            | Package | Uncore M-box 1 perfmon PLD unit select MSR.            |
| CEBH             | 3307 | MSR_M1_PMON_ZDP            | Package | Uncore M-box 1 perfmon ZDP unit select MSR.            |
| CF0H             | 3312 | MSR_M1_PMON_EVNT_SEL0      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CF1H             | 3313 | MSR_M1_PMON_CTR0           | Package | Uncore M-box 1 perfmon counter MSR.                    |
| CF2H             | 3314 | MSR_M1_PMON_EVNT_SEL1      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CF3H             | 3315 | MSR_M1_PMON_CTR1           | Package | Uncore M-box 1 perfmon counter MSR.                    |
| CF4H             | 3316 | MSR_M1_PMON_EVNT_SEL2      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CF5H             | 3317 | MSR_M1_PMON_CTR2           | Package | Uncore M-box 1 perfmon counter MSR.                    |
| CF6H             | 3318 | MSR_M1_PMON_EVNT_SEL3      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CF7H             | 3319 | MSR_M1_PMON_CTR3           | Package | Uncore M-box 1 perfmon counter MSR.                    |
| CF8H             | 3320 | MSR_M1_PMON_EVNT_SEL4      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CF9H             | 3321 | MSR_M1_PMON_CTR4           | Package | Uncore M-box 1 perfmon counter MSR.                    |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| CFAH             | 3322 | MSR_M1_PMON_EVNT_SEL5      | Package | Uncore M-box 1 perfmon event select MSR.               |
| CFBH             | 3323 | MSR_M1_PMON_CTR5           | Package | Uncore M-box 1 perfmon counter MSR.                    |
| DOOH             | 3328 | MSR_C0_PMON_BOX_CTRL       | Package | Uncore C-box 0 perfmon local box control MSR.          |
| D01H             | 3329 | MSR_C0_PMON_BOX_STATUS     | Package | Uncore C-box 0 perfmon local box status MSR.           |
| D02H             | 3330 | MSR_C0_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 0 perfmon local box overflow control MSR. |
| D10H             | 3344 | MSR_C0_PMON_EVNT_SELO      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D11H             | 3345 | MSR_C0_PMON_CTR0           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D12H             | 3346 | MSR_C0_PMON_EVNT_SEL1      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D13H             | 3347 | MSR_C0_PMON_CTR1           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D14H             | 3348 | MSR_C0_PMON_EVNT_SEL2      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D15H             | 3349 | MSR_C0_PMON_CTR2           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D16H             | 3350 | MSR_C0_PMON_EVNT_SEL3      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D17H             | 3351 | MSR_C0_PMON_CTR3           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D18H             | 3352 | MSR_C0_PMON_EVNT_SEL4      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D19H             | 3353 | MSR_C0_PMON_CTR4           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D1AH             | 3354 | MSR_C0_PMON_EVNT_SEL5      | Package | Uncore C-box 0 perfmon event select MSR.               |
| D1BH             | 3355 | MSR_C0_PMON_CTR5           | Package | Uncore C-box 0 perfmon counter MSR.                    |
| D20H             | 3360 | MSR_C4_PMON_BOX_CTRL       | Package | Uncore C-box 4 perfmon local box control MSR.          |
| D21H             | 3361 | MSR_C4_PMON_BOX_STATUS     | Package | Uncore C-box 4 perfmon local box status MSR.           |
| D22H             | 3362 | MSR_C4_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 4 perfmon local box overflow control MSR. |
| D30H             | 3376 | MSR_C4_PMON_EVNT_SELO      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D31H             | 3377 | MSR_C4_PMON_CTR0           | Package | Uncore C-box 4 perfmon counter MSR.                    |
| D32H             | 3378 | MSR_C4_PMON_EVNT_SEL1      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D33H             | 3379 | MSR_C4_PMON_CTR1           | Package | Uncore C-box 4 perfmon counter MSR.                    |
| D34H             | 3380 | MSR_C4_PMON_EVNT_SEL2      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D35H             | 3381 | MSR_C4_PMON_CTR2           | Package | Uncore C-box 4 perfmon counter MSR.                    |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| D36H             | 3382 | MSR_C4_PMON_EVNT_SEL3      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D37H             | 3383 | MSR_C4_PMON_CTR3           | Package | Uncore C-box 4 perfmon counter MSR.                    |
| D38H             | 3384 | MSR_C4_PMON_EVNT_SEL4      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D39H             | 3385 | MSR_C4_PMON_CTR4           | Package | Uncore C-box 4 perfmon counter MSR.                    |
| D3AH             | 3386 | MSR_C4_PMON_EVNT_SEL5      | Package | Uncore C-box 4 perfmon event select MSR.               |
| D3BH             | 3387 | MSR_C4_PMON_CTR5           | Package | Uncore C-box 4 perfmon counter MSR.                    |
| D40H             | 3392 | MSR_C2_PMON_BOX_CTRL       | Package | Uncore C-box 2 perfmon local box control MSR.          |
| D41H             | 3393 | MSR_C2_PMON_BOX_STATUS     | Package | Uncore C-box 2 perfmon local box status MSR.           |
| D42H             | 3394 | MSR_C2_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 2 perfmon local box overflow control MSR. |
| D50H             | 3408 | MSR_C2_PMON_EVNT_SELO      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D51H             | 3409 | MSR_C2_PMON_CTR0           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D52H             | 3410 | MSR_C2_PMON_EVNT_SEL1      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D53H             | 3411 | MSR_C2_PMON_CTR1           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D54H             | 3412 | MSR_C2_PMON_EVNT_SEL2      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D55H             | 3413 | MSR_C2_PMON_CTR2           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D56H             | 3414 | MSR_C2_PMON_EVNT_SEL3      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D57H             | 3415 | MSR_C2_PMON_CTR3           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D58H             | 3416 | MSR_C2_PMON_EVNT_SEL4      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D59H             | 3417 | MSR_C2_PMON_CTR4           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D5AH             | 3418 | MSR_C2_PMON_EVNT_SEL5      | Package | Uncore C-box 2 perfmon event select MSR.               |
| D5BH             | 3419 | MSR_C2_PMON_CTR5           | Package | Uncore C-box 2 perfmon counter MSR.                    |
| D60H             | 3424 | MSR_C6_PMON_BOX_CTRL       | Package | Uncore C-box 6 perfmon local box control MSR.          |
| D61H             | 3425 | MSR_C6_PMON_BOX_STATUS     | Package | Uncore C-box 6 perfmon local box status MSR.           |
| D62H             | 3426 | MSR_C6_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 6 perfmon local box overflow control MSR. |
| D70H             | 3440 | MSR_C6_PMON_EVNT_SELO      | Package | Uncore C-box 6 perfmon event select MSR.               |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| D71H             | 3441 | MSR_C6_PMON_CTRL0          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D72H             | 3442 | MSR_C6_PMON_EVNT_SEL1      | Package | Uncore C-box 6 perfmon event select MSR.               |
| D73H             | 3443 | MSR_C6_PMON_CTRL1          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D74H             | 3444 | MSR_C6_PMON_EVNT_SEL2      | Package | Uncore C-box 6 perfmon event select MSR.               |
| D75H             | 3445 | MSR_C6_PMON_CTRL2          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D76H             | 3446 | MSR_C6_PMON_EVNT_SEL3      | Package | Uncore C-box 6 perfmon event select MSR.               |
| D77H             | 3447 | MSR_C6_PMON_CTRL3          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D78H             | 3448 | MSR_C6_PMON_EVNT_SEL4      | Package | Uncore C-box 6 perfmon event select MSR.               |
| D79H             | 3449 | MSR_C6_PMON_CTRL4          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D7AH             | 3450 | MSR_C6_PMON_EVNT_SEL5      | Package | Uncore C-box 6 perfmon event select MSR.               |
| D7BH             | 3451 | MSR_C6_PMON_CTRL5          | Package | Uncore C-box 6 perfmon counter MSR.                    |
| D80H             | 3456 | MSR_C1_PMON_BOX_CTRL       | Package | Uncore C-box 1 perfmon local box control MSR.          |
| D81H             | 3457 | MSR_C1_PMON_BOX_STATUS     | Package | Uncore C-box 1 perfmon local box status MSR.           |
| D82H             | 3458 | MSR_C1_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 1 perfmon local box overflow control MSR. |
| D90H             | 3472 | MSR_C1_PMON_EVNT_SEL0      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D91H             | 3473 | MSR_C1_PMON_CTRL0          | Package | Uncore C-box 1 perfmon counter MSR.                    |
| D92H             | 3474 | MSR_C1_PMON_EVNT_SEL1      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D93H             | 3475 | MSR_C1_PMON_CTRL1          | Package | Uncore C-box 1 perfmon counter MSR.                    |
| D94H             | 3476 | MSR_C1_PMON_EVNT_SEL2      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D95H             | 3477 | MSR_C1_PMON_CTRL2          | Package | Uncore C-box 1 perfmon counter MSR.                    |
| D96H             | 3478 | MSR_C1_PMON_EVNT_SEL3      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D97H             | 3479 | MSR_C1_PMON_CTRL3          | Package | Uncore C-box 1 perfmon counter MSR.                    |
| D98H             | 3480 | MSR_C1_PMON_EVNT_SEL4      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D99H             | 3481 | MSR_C1_PMON_CTRL4          | Package | Uncore C-box 1 perfmon counter MSR.                    |
| D9AH             | 3482 | MSR_C1_PMON_EVNT_SEL5      | Package | Uncore C-box 1 perfmon event select MSR.               |
| D9BH             | 3483 | MSR_C1_PMON_CTRL5          | Package | Uncore C-box 1 perfmon counter MSR.                    |



Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| DA0H             | 3488 | MSR_C5_PMON_BOX_CTRL       | Package | Uncore C-box 5 perfmon local box control MSR.          |
| DA1H             | 3489 | MSR_C5_PMON_BOX_STATUS     | Package | Uncore C-box 5 perfmon local box status MSR.           |
| DA2H             | 3490 | MSR_C5_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 5 perfmon local box overflow control MSR. |
| DB0H             | 3504 | MSR_C5_PMON_EVNT_SELO      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DB1H             | 3505 | MSR_C5_PMON_CTR0           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DB2H             | 3506 | MSR_C5_PMON_EVNT_SEL1      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DB3H             | 3507 | MSR_C5_PMON_CTR1           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DB4H             | 3508 | MSR_C5_PMON_EVNT_SEL2      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DB5H             | 3509 | MSR_C5_PMON_CTR2           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DB6H             | 3510 | MSR_C5_PMON_EVNT_SEL3      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DB7H             | 3511 | MSR_C5_PMON_CTR3           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DB8H             | 3512 | MSR_C5_PMON_EVNT_SEL4      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DB9H             | 3513 | MSR_C5_PMON_CTR4           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DBAH             | 3514 | MSR_C5_PMON_EVNT_SEL5      | Package | Uncore C-box 5 perfmon event select MSR.               |
| DBBH             | 3515 | MSR_C5_PMON_CTR5           | Package | Uncore C-box 5 perfmon counter MSR.                    |
| DC0H             | 3520 | MSR_C3_PMON_BOX_CTRL       | Package | Uncore C-box 3 perfmon local box control MSR.          |
| DC1H             | 3521 | MSR_C3_PMON_BOX_STATUS     | Package | Uncore C-box 3 perfmon local box status MSR.           |
| DC2H             | 3522 | MSR_C3_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 3 perfmon local box overflow control MSR. |
| DD0H             | 3536 | MSR_C3_PMON_EVNT_SELO      | Package | Uncore C-box 3 perfmon event select MSR.               |
| DD1H             | 3537 | MSR_C3_PMON_CTR0           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DD2H             | 3538 | MSR_C3_PMON_EVNT_SEL1      | Package | Uncore C-box 3 perfmon event select MSR.               |
| DD3H             | 3539 | MSR_C3_PMON_CTR1           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DD4H             | 3540 | MSR_C3_PMON_EVNT_SEL2      | Package | Uncore C-box 3 perfmon event select MSR.               |
| DD5H             | 3541 | MSR_C3_PMON_CTR2           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DD6H             | 3542 | MSR_C3_PMON_EVNT_SEL3      | Package | Uncore C-box 3 perfmon event select MSR.               |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| DD7H             | 3543 | MSR_C3_PMON_CTR3           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DD8H             | 3544 | MSR_C3_PMON_EVNT_SEL4      | Package | Uncore C-box 3 perfmon event select MSR.               |
| DD9H             | 3545 | MSR_C3_PMON_CTR4           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DDAH             | 3546 | MSR_C3_PMON_EVNT_SEL5      | Package | Uncore C-box 3 perfmon event select MSR.               |
| DDBH             | 3547 | MSR_C3_PMON_CTR5           | Package | Uncore C-box 3 perfmon counter MSR.                    |
| DE0H             | 3552 | MSR_C7_PMON_BOX_CTRL       | Package | Uncore C-box 7 perfmon local box control MSR.          |
| DE1H             | 3553 | MSR_C7_PMON_BOX_STATUS     | Package | Uncore C-box 7 perfmon local box status MSR.           |
| DE2H             | 3554 | MSR_C7_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 7 perfmon local box overflow control MSR. |
| DF0H             | 3568 | MSR_C7_PMON_EVNT_SELO      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DF1H             | 3569 | MSR_C7_PMON_CTR0           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| DF2H             | 3570 | MSR_C7_PMON_EVNT_SEL1      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DF3H             | 3571 | MSR_C7_PMON_CTR1           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| DF4H             | 3572 | MSR_C7_PMON_EVNT_SEL2      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DF5H             | 3573 | MSR_C7_PMON_CTR2           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| DF6H             | 3574 | MSR_C7_PMON_EVNT_SEL3      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DF7H             | 3575 | MSR_C7_PMON_CTR3           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| DF8H             | 3576 | MSR_C7_PMON_EVNT_SEL4      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DF9H             | 3577 | MSR_C7_PMON_CTR4           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| DFAH             | 3578 | MSR_C7_PMON_EVNT_SEL5      | Package | Uncore C-box 7 perfmon event select MSR.               |
| DFBH             | 3579 | MSR_C7_PMON_CTR5           | Package | Uncore C-box 7 perfmon counter MSR.                    |
| E00H             | 3584 | MSR_R0_PMON_BOX_CTRL       | Package | Uncore R-box 0 perfmon local box control MSR.          |
| E01H             | 3585 | MSR_R0_PMON_BOX_STATUS     | Package | Uncore R-box 0 perfmon local box status MSR.           |
| E02H             | 3586 | MSR_R0_PMON_BOX_OVF_CTRL   | Package | Uncore R-box 0 perfmon local box overflow control MSR. |
| E04H             | 3588 | MSR_R0_PMON_IPERFO_P0      | Package | Uncore R-box 0 perfmon IPERFO unit Port 0 select MSR.  |
| E05H             | 3589 | MSR_R0_PMON_IPERFO_P1      | Package | Uncore R-box 0 perfmon IPERFO unit Port 1 select MSR.  |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description                                       |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| E06H             | 3590 | MSR_R0_PMON_IPERF0_P2      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 2 select MSR. |
| E07H             | 3591 | MSR_R0_PMON_IPERF0_P3      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 3 select MSR. |
| E08H             | 3592 | MSR_R0_PMON_IPERF0_P4      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 4 select MSR. |
| E09H             | 3593 | MSR_R0_PMON_IPERF0_P5      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 5 select MSR. |
| E0AH             | 3594 | MSR_R0_PMON_IPERF0_P6      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 6 select MSR. |
| E0BH             | 3595 | MSR_R0_PMON_IPERF0_P7      | Package | Uncore R-box 0 perfmon IPERF0 unit Port 7 select MSR. |
| E0CH             | 3596 | MSR_R0_PMON_QLX_P0         | Package | Uncore R-box 0 perfmon QLX unit Port 0 select MSR.    |
| E0DH             | 3597 | MSR_R0_PMON_QLX_P1         | Package | Uncore R-box 0 perfmon QLX unit Port 1 select MSR.    |
| E0EH             | 3598 | MSR_R0_PMON_QLX_P2         | Package | Uncore R-box 0 perfmon QLX unit Port 2 select MSR.    |
| E0FH             | 3599 | MSR_R0_PMON_QLX_P3         | Package | Uncore R-box 0 perfmon QLX unit Port 3 select MSR.    |
| E10H             | 3600 | MSR_R0_PMON_EVNT_SEL0      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E11H             | 3601 | MSR_R0_PMON_CTR0           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E12H             | 3602 | MSR_R0_PMON_EVNT_SEL1      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E13H             | 3603 | MSR_R0_PMON_CTR1           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E14H             | 3604 | MSR_R0_PMON_EVNT_SEL2      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E15H             | 3605 | MSR_R0_PMON_CTR2           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E16H             | 3606 | MSR_R0_PMON_EVNT_SEL3      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E17H             | 3607 | MSR_R0_PMON_CTR3           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E18H             | 3608 | MSR_R0_PMON_EVNT_SEL4      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E19H             | 3609 | MSR_R0_PMON_CTR4           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E1AH             | 3610 | MSR_R0_PMON_EVNT_SEL5      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E1BH             | 3611 | MSR_R0_PMON_CTR5           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E1CH             | 3612 | MSR_R0_PMON_EVNT_SEL6      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E1DH             | 3613 | MSR_R0_PMON_CTR6           | Package | Uncore R-box 0 perfmon counter MSR.                   |
| E1EH             | 3614 | MSR_R0_PMON_EVNT_SEL7      | Package | Uncore R-box 0 perfmon event select MSR.              |
| E1FH             | 3615 | MSR_R0_PMON_CTR7           | Package | Uncore R-box 0 perfmon counter MSR.                   |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| E20H             | 3616 | MSR_R1_PMON_BOX_CTRL       | Package | Uncore R-box 1 perfmon local box control MSR.          |
| E21H             | 3617 | MSR_R1_PMON_BOX_STATUS     | Package | Uncore R-box 1 perfmon local box status MSR.           |
| E22H             | 3618 | MSR_R1_PMON_BOX_OVF_CTRL   | Package | Uncore R-box 1 perfmon local box overflow control MSR. |
| E24H             | 3620 | MSR_R1_PMON_IPERF1_P8      | Package | Uncore R-box 1 perfmon IPERF1 unit Port 8 select MSR.  |
| E25H             | 3621 | MSR_R1_PMON_IPERF1_P9      | Package | Uncore R-box 1 perfmon IPERF1 unit Port 9 select MSR.  |
| E26H             | 3622 | MSR_R1_PMON_IPERF1_P10     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 10 select MSR. |
| E27H             | 3623 | MSR_R1_PMON_IPERF1_P11     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 11 select MSR. |
| E28H             | 3624 | MSR_R1_PMON_IPERF1_P12     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 12 select MSR. |
| E29H             | 3625 | MSR_R1_PMON_IPERF1_P13     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 13 select MSR. |
| E2AH             | 3626 | MSR_R1_PMON_IPERF1_P14     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 14 select MSR. |
| E2BH             | 3627 | MSR_R1_PMON_IPERF1_P15     | Package | Uncore R-box 1 perfmon IPERF1 unit Port 15 select MSR. |
| E2CH             | 3628 | MSR_R1_PMON_QLX_P4         | Package | Uncore R-box 1 perfmon QLX unit Port 4 select MSR.     |
| E2DH             | 3629 | MSR_R1_PMON_QLX_P5         | Package | Uncore R-box 1 perfmon QLX unit Port 5 select MSR.     |
| E2EH             | 3630 | MSR_R1_PMON_QLX_P6         | Package | Uncore R-box 1 perfmon QLX unit Port 6 select MSR.     |
| E2FH             | 3631 | MSR_R1_PMON_QLX_P7         | Package | Uncore R-box 1 perfmon QLX unit Port 7 select MSR.     |
| E30H             | 3632 | MSR_R1_PMON_EVNT_SEL8      | Package | Uncore R-box 1 perfmon event select MSR.               |
| E31H             | 3633 | MSR_R1_PMON_CTR8           | Package | Uncore R-box 1 perfmon counter MSR.                    |
| E32H             | 3634 | MSR_R1_PMON_EVNT_SEL9      | Package | Uncore R-box 1 perfmon event select MSR.               |
| E33H             | 3635 | MSR_R1_PMON_CTR9           | Package | Uncore R-box 1 perfmon counter MSR.                    |
| E34H             | 3636 | MSR_R1_PMON_EVNT_SEL10     | Package | Uncore R-box 1 perfmon event select MSR.               |
| E35H             | 3637 | MSR_R1_PMON_CTR10          | Package | Uncore R-box 1 perfmon counter MSR.                    |
| E36H             | 3638 | MSR_R1_PMON_EVNT_SEL11     | Package | Uncore R-box 1 perfmon event select MSR.               |
| E37H             | 3639 | MSR_R1_PMON_CTR11          | Package | Uncore R-box 1 perfmon counter MSR.                    |
| E38H             | 3640 | MSR_R1_PMON_EVNT_SEL12     | Package | Uncore R-box 1 perfmon event select MSR.               |
| E39H             | 3641 | MSR_R1_PMON_CTR12          | Package | Uncore R-box 1 perfmon counter MSR.                    |

Table 2-16. Additional MSRs in Intel® Xeon® Processor 7500 Series (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| E3AH             | 3642 | MSR_R1_PMON_EVNT_SEL13     | Package | Uncore R-box 1 perfmon event select MSR.                                   |
| E3BH             | 3643 | MSR_R1_PMON_CTR13          | Package | Uncore R-box 1perfmon counter MSR.   |
| E3CH             | 3644 | MSR_R1_PMON_EVNT_SEL14     | Package | Uncore R-box 1 perfmon event select MSR.                                   |
| E3DH             | 3645 | MSR_R1_PMON_CTR14          | Package | Uncore R-box 1 perfmon counter MSR.  |
| E3EH             | 3646 | MSR_R1_PMON_EVNT_SEL15     | Package | Uncore R-box 1 perfmon event select MSR.                                   |
| E3FH             | 3647 | MSR_R1_PMON_CTR15          | Package | Uncore R-box 1 perfmon counter MSR.  |
| E45H             | 3653 | MSR_B0_PMON_MATCH          | Package | Uncore B-box 0 perfmon local box match MSR.                                |
| E46H             | 3654 | MSR_B0_PMON_MASK           | Package | Uncore B-box 0 perfmon local box mask MSR.                                 |
| E49H             | 3657 | MSR_S0_PMON_MATCH          | Package | Uncore S-box 0 perfmon local box match MSR.                                |
| E4AH             | 3658 | MSR_S0_PMON_MASK           | Package | Uncore S-box 0 perfmon local box mask MSR.                                 |
| E4DH             | 3661 | MSR_B1_PMON_MATCH          | Package | Uncore B-box 1 perfmon local box match MSR.                                |
| E4EH             | 3662 | MSR_B1_PMON_MASK           | Package | Uncore B-box 1 perfmon local box mask MSR.                                 |
| E54H             | 3668 | MSR_M0_PMON_MM_CONFIG      | Package | Uncore M-box 0 perfmon local box address match/mask config MSR.            |
| E55H             | 3669 | MSR_M0_PMON_ADDR_MATCH     | Package | Uncore M-box 0 perfmon local box address match MSR.                        |
| E56H             | 3670 | MSR_M0_PMON_ADDR_MASK      | Package | Uncore M-box 0 perfmon local box address mask MSR.                         |
| E59H             | 3673 | MSR_S1_PMON_MATCH          | Package | Uncore S-box 1 perfmon local box match MSR.                                |
| E5AH             | 3674 | MSR_S1_PMON_MASK           | Package | Uncore S-box 1 perfmon local box mask MSR.                                 |
| E5CH             | 3676 | MSR_M1_PMON_MM_CONFIG      | Package | Uncore M-box 1 perfmon local box address match/mask config MSR.            |
| E5DH             | 3677 | MSR_M1_PMON_ADDR_MATCH     | Package | Uncore M-box 1 perfmon local box address match MSR.                        |
| E5EH             | 3678 | MSR_M1_PMON_ADDR_MASK      | Package | Uncore M-box 1 perfmon local box address mask MSR.                         |
| 3B5H             | 965  | MSR_UNCORE_PMC5            | Package | See Section 18.3.1.2.2, "Uncore Performance Event Configuration Facility." |

## 2.8 MSRS IN THE INTEL® XEON® PROCESSOR 5600 SERIES (BASED ON INTEL® MICROARCHITECTURE CODE NAME WESTMERE)

Intel® Xeon® Processor 5600 Series (based on Intel® microarchitecture code name Westmere) supports the MSR interfaces listed in Table 2-14, Table 2-15, plus additional MSR listed in Table 2-17. These MSRs apply to Intel Core i7, i5 and i3 processor family with CPUID signature DisplayFamily\_DisplayModel of 06\_25H and 06\_2CH, see Table 2-1.

**Table 2-17. Additional MSRs Supported by Intel Processors  
(Based on Intel® Microarchitecture Code Name Westmere)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 13CH             | 52  | MSR_FEATURE_CONFIG         | Core    | AES Configuration (RW-L)<br>Privileged post-BIOS agent must provide a #GP handler to handle unsuccessful read of this MSR.   |
|                  |     | 1:0                        |         | AES Configuration (RW-L)<br>Upon a successful read of this MSR, the configuration of AES instruction set availability is as follows:<br>11b: AES instructions are not available until next RESET.<br>Otherwise, AES instructions are available.<br>Note, AES instruction set is not available if read is unsuccessful. If the configuration is not 01b, AES instructions can be mis-configured if a privileged agent unintentionally writes 11b. |
|                  |     | 63:2                       |         | Reserved   |
| 1A7H             | 423 | MSR_OFFCORE_RSP_1          | Thread  | Offcore Response Event Select Register (R/W)   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |
|                  |     | 39:32                      | Package | Maximum Ratio Limit for 5C<br>Maximum turbo ratio limit of 5 core active.  |
|                  |     | 47:40                      | Package | Maximum Ratio Limit for 6C<br>Maximum turbo ratio limit of 6 core active.  |
|                  |     | 63:48                      |         | Reserved   |
| 1B0H             | 432 | IA32_ENERGY_PERF_BIAS      | Package | See Table 2-2.   |

## 2.9 MSRS IN THE INTEL® XEON® PROCESSOR E7 FAMILY (BASED ON INTEL® MICROARCHITECTURE CODE NAME WESTMERE)

Intel® Xeon® Processor E7 Family (based on Intel® microarchitecture code name Westmere) supports the MSR interfaces listed in Table 2-14 (except MSR address 1ADH), Table 2-15, plus additional MSR listed in Table 2-18. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2FH.

Table 2-18. Additional MSRs Supported by Intel® Xeon® Processor E7 Family

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 13CH             | 52   | MSR_FEATURE_CONFIG         | Core    | AES Configuration (RW-L)<br>Privileged post-BIOS agent must provide a #GP handler to handle unsuccessful read of this MSR.  |
|                  |      | 1:0                        |         | AES Configuration (RW-L)<br>Upon a successful read of this MSR, the configuration of AES instruction set availability is as follows:<br>11b: AES instructions are not available until next RESET. Otherwise, AES instructions are available.<br>Note, AES instruction set is not available if read is unsuccessful. If the configuration is not 01b, AES instructions can be mis-configured if a privileged agent unintentionally writes 11b. |
|                  |      | 63:2                       |         | Reserved  |
| 1A7H             | 423  | MSR_OFFCORE_RSP_1          | Thread  | Offcore Response Event Select Register (R/W)  |
| 1ADH             | 429  | MSR_TURBO_RATIO_LIMIT      | Package | Reserved<br>Attempt to read/write will cause #UD.   |
| 1B0H             | 432  | IA32_ENERGY_PERF_BIAS      | Package | See Table 2-2.  |
| F40H             | 3904 | MSR_C8_PMON_BOX_CTRL       | Package | Uncore C-box 8 perfmon local box control MSR.   |
| F41H             | 3905 | MSR_C8_PMON_BOX_STATUS     | Package | Uncore C-box 8 perfmon local box status MSR.  |
| F42H             | 3906 | MSR_C8_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 8 perfmon local box overflow control MSR.  |
| F50H             | 3920 | MSR_C8_PMON_EVNT_SELO      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F51H             | 3921 | MSR_C8_PMON_CTR0           | Package | Uncore C-box 8 perfmon counter MSR.   |
| F52H             | 3922 | MSR_C8_PMON_EVNT_SEL1      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F53H             | 3923 | MSR_C8_PMON_CTR1           | Package | Uncore C-box 8 perfmon counter MSR.   |
| F54H             | 3924 | MSR_C8_PMON_EVNT_SEL2      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F55H             | 3925 | MSR_C8_PMON_CTR2           | Package | Uncore C-box 8 perfmon counter MSR.   |
| F56H             | 3926 | MSR_C8_PMON_EVNT_SEL3      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F57H             | 3927 | MSR_C8_PMON_CTR3           | Package | Uncore C-box 8 perfmon counter MSR.   |
| F58H             | 3928 | MSR_C8_PMON_EVNT_SEL4      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F59H             | 3929 | MSR_C8_PMON_CTR4           | Package | Uncore C-box 8 perfmon counter MSR.   |
| F5AH             | 3930 | MSR_C8_PMON_EVNT_SEL5      | Package | Uncore C-box 8 perfmon event select MSR.  |
| F5BH             | 3931 | MSR_C8_PMON_CTR5           | Package | Uncore C-box 8 perfmon counter MSR.   |

**Table 2-18. Additional MSRs Supported by Intel® Xeon® Processor E7 Family (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| FC0H             | 4032 | MSR_C9_PMON_BOX_CTRL       | Package | Uncore C-box 9 perfmon local box control MSR.          |
| FC1H             | 4033 | MSR_C9_PMON_BOX_STATUS     | Package | Uncore C-box 9 perfmon local box status MSR.           |
| FC2H             | 4034 | MSR_C9_PMON_BOX_OVF_CTRL   | Package | Uncore C-box 9 perfmon local box overflow control MSR. |
| FD0H             | 4048 | MSR_C9_PMON_EVNT_SELO      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FD1H             | 4049 | MSR_C9_PMON_CTRL0          | Package | Uncore C-box 9 perfmon counter MSR.                    |
| FD2H             | 4050 | MSR_C9_PMON_EVNT_SEL1      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FD3H             | 4051 | MSR_C9_PMON_CTRL1          | Package | Uncore C-box 9 perfmon counter MSR.                    |
| FD4H             | 4052 | MSR_C9_PMON_EVNT_SEL2      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FD5H             | 4053 | MSR_C9_PMON_CTRL2          | Package | Uncore C-box 9 perfmon counter MSR.                    |
| FD6H             | 4054 | MSR_C9_PMON_EVNT_SEL3      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FD7H             | 4055 | MSR_C9_PMON_CTRL3          | Package | Uncore C-box 9 perfmon counter MSR.                    |
| FD8H             | 4056 | MSR_C9_PMON_EVNT_SEL4      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FD9H             | 4057 | MSR_C9_PMON_CTRL4          | Package | Uncore C-box 9 perfmon counter MSR.                    |
| FDAH             | 4058 | MSR_C9_PMON_EVNT_SEL5      | Package | Uncore C-box 9 perfmon event select MSR.               |
| FDBH             | 4059 | MSR_C9_PMON_CTRL5          | Package | Uncore C-box 9 perfmon counter MSR.                    |

## 2.10 MSRS IN INTEL® PROCESSOR FAMILY BASED ON INTEL® MICROARCHITECTURE CODE NAME SANDY BRIDGE

Table 2-19 lists model-specific registers (MSRs) that are common to Intel® processor family based on Intel microarchitecture code name Sandy Bridge. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2AH, 06\_2DH, see Table 2-1. Additional MSRs specific to 06\_2AH are listed in Table 2-20.

**Table 2-19. MSRs Supported by Intel® Processors based on Intel® microarchitecture code name Sandy Bridge**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 0H               | 0   | IA32_P5_MC_ADDR            | Thread | See Section 2.22, “MSRs in Pentium Processors.”                                |
| 1H               | 1   | IA32_P5_MC_TYPE            | Thread | See Section 2.22, “MSRs in Pentium Processors.”                                |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Thread | See Section 8.10.5, “Monitor/Mwait Address Range Determination” and Table 2-2. |



**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope                                 | Bit Description   |
|------------------|-----|----------------------------|---------------------------------------|---|
| Hex              | Dec |                            |                                       |   |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Thread                                | See Section 17.17, "Time-Stamp Counter" and see Table 2-2.          |
| 17H              | 23  | IA32_PLATFORM_ID           | Package                               | Platform ID (R)<br>See Table 2-2.                                   |
| 1BH              | 27  | IA32_APIC_BASE             | Thread                                | See Section 10.4.4, "Local APIC Status and Location" and Table 2-2. |
| 34H              | 52  | MSR_SMI_COUNT              | Thread                                | SMI Counter (R/O)   |
|                  |     | 31:0                       |                                       | SMI Count (R/O)<br>Count SMIs.                                      |
|                  |     | 63:32                      |                                       | Reserved.   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread                                | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2.      |
|                  |     | 0                          |                                       | Lock (R/WL)   |
|                  |     | 1                          |                                       | Enable VMX Inside SMX Operation (R/WL)                              |
|                  |     | 2                          |                                       | Enable VMX Outside SMX Operation (R/WL)                             |
|                  |     | 14:8                       |                                       | SENTER Local Functions Enables (R/WL)                               |
|                  | 15  |                            | SENTER Global Functions Enable (R/WL) |   |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Core                                  | BIOS Update Trigger Register (W)<br>See Table 2-2.                  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | Thread                                | BIOS Update Signature ID (RO)<br>See Table 2-2.                     |
| C1H              | 193 | IA32_PMC0                  | Thread                                | Performance Counter Register<br>See Table 2-2.                      |
| C2H              | 194 | IA32_PMC1                  | Thread                                | Performance Counter Register<br>See Table 2-2.                      |
| C3H              | 195 | IA32_PMC2                  | Thread                                | Performance Counter Register<br>See Table 2-2.                      |
| C4H              | 196 | IA32_PMC3                  | Thread                                | Performance Counter Register<br>See Table 2-2.                      |
| C5H              | 197 | IA32_PMC4                  | Core                                  | Performance Counter Register (if core not shared by threads)        |
| C6H              | 198 | IA32_PMC5                  | Core                                  | Performance Counter Register (if core not shared by threads)        |
| C7H              | 199 | IA32_PMC6                  | Core                                  | Performance Counter Register (if core not shared by threads)        |
| C8H              | 200 | IA32_PMC7                  | Core                                  | Performance Counter Register (if core not shared by threads)        |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the frequency that invariant TSC runs at. Frequency = ratio * 100 MHz.   |
|                  |     | 27:16                      |         | Reserved   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled.   |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates TDP Limit for Turbo mode is not programmable.  |
|                  |     | 39:30                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz.   |
|                  |     | 63:48                      |         | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br><br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-sate support)<br>001b: C2<br>010b: C6 no retention<br>011b: C6 retention<br>100b: C7<br>101b: C7s<br>111: No package C-state limit<br>Note: This field cannot be used to limit package C-state to C3. |
|                  |     | 9:3                        |         | Reserved   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 10                         |        | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.  |
|                  |     | 14:11                      |        | Reserved   |
|                  |     | 15                         |        | CFG Lock (R/W0)<br>When set, locks bits 15:0 of this register until next reset.  |
|                  |     | 24:16                      |        | Reserved   |
|                  |     | 25                         |        | C3 State Auto Demotion Enable (R/W)<br>When set, the processor will conditionally demote C6/C7 requests to C3 based on uncore auto-demote information.   |
|                  |     | 26                         |        | C1 State Auto Demotion Enable (R/W)<br>When set, the processor will conditionally demote C3/C6/C7 requests to C1 based on uncore auto-demote information.  |
|                  |     | 27                         |        | Enable C3 Undemotion (R/W)<br>When set, enables undemotion from demoted C3.  |
|                  |     | 28                         |        | Enable C1 Undemotion (R/W)<br>When set, enables undemotion from demoted C1.  |
|                  |     | 63:29                      |        | Reserved   |
| E4H              | 228 | MSR_PMG_IO_CAPTURE_BASE    | Core   | Power Management IO Redirection in C-state (R/W)<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 15:0                       |        | LVL_2 Base Address (R/W)<br>Specifies the base address visible to software for IO redirection. If IO MWAIT Redirection is enabled, reads to this address will be consumed by the power management logic and decoded to MWAIT instructions. When IO port address redirection is enabled, this is the IO port address reported to the OS/software. |
|                  |     | 18:16                      |        | C-State Range (R/W)<br>Specifies the encoding value of the maximum C-State code name to be included when IO read to MWAIT redirection is enabled by MSR_PKG_CST_CONFIG_CONTROL[bit10]:<br>000b - C3 is the max C-State to include.<br>001b - C6 is the max C-State to include.<br>010b - C7 is the max C-State to include.                       |
|                  |     | 63:19                      |        | Reserved   |
| E7H              | 231 | IA32_MPERF                 | Thread | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| E8H              | 232 | IA32_APERF                 | Thread | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP               | Thread | See Table 2-2.   |
| 13CH             | 52  | MSR_FEATURE_CONFIG         | Core   | AES Configuration (RW-L)<br>Privileged post-BIOS agent must provide a #GP handler to handle unsuccessful read of this MSR.   |
|                  |     | 1:0                        |        | AES Configuration (RW-L)<br>Upon a successful read of this MSR, the configuration of AES instruction set availability is as follows:<br>11b: AES instructions are not available until next RESET.<br>Otherwise, AES instructions are available.<br>Note, AES instruction set is not available if read is unsuccessful. If the configuration is not 01b, AES instructions can be mis-configured if a privileged agent unintentionally writes 11b. |
|                  |     | 63:2                       |        | Reserved   |
| 174H             | 372 | IA32_SYSENTER_CS           | Thread | See Table 2-2.   |
| 175H             | 373 | IA32_SYSENTER_ESP          | Thread | See Table 2-2.   |
| 176H             | 374 | IA32_SYSENTER_EIP          | Thread | See Table 2-2.   |
| 179H             | 377 | IA32_MCG_CAP               | Thread | See Table 2-2.   |
| 17AH             | 378 | IA32_MCG_STATUS            | Thread | Global Machine Check Status  |
|                  |     | 0                          |        | RIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If cleared, the program cannot be reliably restarted.   |
|                  |     | 1                          |        | EIPV<br>When set, bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |
|                  |     | 2                          |        | MCIP<br>When set, bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception.   |
| 63:3             |     | Reserved                   |        |  |
| 186H             | 390 | IA32_PERFEVTSELO           | Thread | See Table 2-2.   |
| 187H             | 391 | IA32_PERFEVTSEL1           | Thread | See Table 2-2.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 188H             | 392 | IA32_PERFEVTSEL2           | Thread  | See Table 2-2.  |
| 189H             | 393 | IA32_PERFEVTSEL3           | Thread  | See Table 2-2.  |
| 18AH             | 394 | IA32_PERFEVTSEL4           | Core    | See Table 2-2. If CPUID.0AH:EAX[15:8] = 8.  |
| 18BH             | 395 | IA32_PERFEVTSEL5           | Core    | See Table 2-2. If CPUID.0AH:EAX[15:8] = 8.  |
| 18CH             | 396 | IA32_PERFEVTSEL6           | Core    | See Table 2-2. If CPUID.0AH:EAX[15:8] = 8.  |
| 18DH             | 397 | IA32_PERFEVTSEL7           | Core    | See Table 2-2. If CPUID.0AH:EAX[15:8] = 8.  |
| 198H             | 408 | IA32_PERF_STATUS           | Package | See Table 2-2.  |
|                  |     | 15:0                       |         | Current Performance State Value   |
|                  |     | 63:16                      |         | Reserved  |
| 198H             | 408 | MSR_PERF_STATUS            | Package | Performance Status  |
|                  |     | 47:32                      |         | Core Voltage (R/O)<br>P-state core voltage can be computed by<br>MSR_PERF_STATUS[37:32] * (float) 1/(2 <sup>13</sup> ). |
| 199H             | 409 | IA32_PERF_CTL              | Thread  | See Table 2-2.  |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Thread  | Clock Modulation (R/W)<br>See Table 2-2.<br>IA32_CLOCK_MODULATION MSR was originally named<br>IA32_THERM_CONTROL MSR.   |
|                  |     | 3:0                        |         | On demand Clock Modulation Duty Cycle (R/W)<br>In 6.25% increment.  |
|                  |     | 4                          |         | On demand Clock Modulation Enable (R/W)   |
|                  |     | 63:5                       |         | Reserved  |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Core    | Thermal Interrupt Control (R/W)<br>See Table 2-2.   |
| 19CH             | 412 | IA32_THERM_STATUS          | Core    | Thermal Monitor Status (R/W)<br>See Table 2-2.  |
|                  |     | 0                          |         | Thermal Status (RO)<br>See Table 2-2.   |
|                  |     | 1                          |         | Thermal Status Log (R/WCO)<br>See Table 2-2.  |
|                  |     | 2                          |         | PROTCHOT # or FORCEPR# Status (RO)<br>See Table 2-2.  |
|                  |     | 3                          |         | PROTCHOT # or FORCEPR# Log (R/WCO)<br>See Table 2-2.  |
|                  |     | 4                          |         | Critical Temperature Status (RO)<br>See Table 2-2.  |
|                  |     | 5                          |         | Critical Temperature Status Log (R/WCO)<br>See Table 2-2.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 6                          |         | Thermal Threshold #1 Status (RO)<br>See Table 2-2.   |
|                  |     | 7                          |         | Thermal Threshold #1 Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 8                          |         | Thermal Threshold #2 Status (RO)<br>See Table 2-2.   |
|                  |     | 9                          |         | Thermal Threshold #2 Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 10                         |         | Power Limitation Status (RO)<br>See Table 2-2.   |
|                  |     | 11                         |         | Power Limitation Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 15:12                      |         | Reserved   |
|                  |     | 22:16                      |         | Digital Readout (RO)<br>See Table 2-2.   |
|                  |     | 26:23                      |         | Reserved   |
|                  |     | 30:27                      |         | Resolution in Degrees Celsius (RO)<br>See Table 2-2.   |
|                  |     | 31                         |         | Reading Valid (RO)<br>See Table 2-2.   |
|                  |     | 63:32                      |         | Reserved   |
| 1A0H             | 416 | IA32_MISC_ENABLE           |         | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled. |
|                  |     | 0                          | Thread  | Fast-Strings Enable<br>See Table 2-2   |
|                  |     | 6:1                        |         | Reserved   |
|                  |     | 7                          | Thread  | Performance Monitoring Available (R)<br>See Table 2-2.   |
|                  |     | 10:8                       |         | Reserved   |
|                  |     | 11                         | Thread  | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.  |
|                  |     | 12                         | Thread  | Processor Event Based Sampling Unavailable (RO)<br>See Table 2-2.  |
|                  |     | 15:13                      |         | Reserved   |
|                  |     | 16                         | Package | Enhanced Intel SpeedStep Technology Enable (R/W)<br>See Table 2-2.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 18                         | Thread  | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.   |
|                  |     | 21:19                      |         | Reserved   |
|                  |     | 22                         | Thread  | Limit CPUID Maxval (R/W)<br>See Table 2-2.   |
|                  |     | 23                         | Thread  | xTPR Message Disable (R/W)<br>See Table 2-2.   |
|                  |     | 33:24                      |         | Reserved   |
|                  |     | 34                         | Thread  | XD Bit Disable (R/W)<br>See Table 2-2.   |
|                  |     | 37:35                      |         | Reserved   |
|                  |     | 38                         | Package | Turbo Mode Disable (R/W)<br>When set to 1 on processors that support Intel Turbo Boost Technology, the turbo mode feature is disabled and the IDA_Enable feature flag will be clear (CPUID.06H: EAX[1]=0).<br>When set to a 0 on processors that support IDA, CPUID.06H: EAX[1] reports the processor's support of turbo mode is enabled.<br>Note: The power-on default value is used by BIOS to detect hardware support of turbo mode. If the power-on default value is 1, turbo mode is available in the processor. If the power-on default value is 0, turbo mode is not available. |
|                  |     | 63:39                      |         | Reserved   |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Unique  | Temperature Target   |
|                  |     | 15:0                       |         | Reserved   |
|                  |     | 23:16                      |         | Temperature Target (R)<br>The minimum temperature at which PROCHOT# will be asserted. The value is degrees C.  |
|                  |     | 63:24                      |         | Reserved   |
| 1A4H             | 420 | MSR_MISC_FEATURE_CONTROL   |         | Miscellaneous Feature Control (R/W)  |
|                  |     | 0                          | Core    | L2 Hardware Prefetcher Disable (R/W)<br>If 1, disables the L2 hardware prefetcher, which fetches additional lines of code or data into the L2 cache.   |
|                  |     | 1                          | Core    | L2 Adjacent Cache Line Prefetcher Disable (R/W)<br>If 1, disables the adjacent cache line prefetcher, which fetches the cache line that comprises a cache line pair (128 bytes).   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields   | Scope   | Bit Description   |
|------------------|-----|------------------------------|---------|---|
| Hex              | Dec |                              |         |   |
|                  |     | 2                            | Core    | DCU Hardware Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache prefetcher, which fetches the next cache line into L1 data cache.   |
|                  |     | 3                            | Core    | DCU IP Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache IP prefetcher, which uses sequential load history (based on instruction pointer of previous loads) to determine whether to prefetch additional lines. |
|                  |     | 63:4                         |         | Reserved  |
| 1A6H             | 422 | MSR_OFFCORE_RSP_0            | Thread  | Offcore Response Event Select Register (R/W)  |
| 1A7H             | 422 | MSR_OFFCORE_RSP_1            | Thread  | Offcore Response Event Select Register (R/W)  |
| 1AAH             | 426 | MSR_MISC_PWR_MGMT            |         | Miscellaneous Power Management Control<br>Various model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .  |
| 1BOH             | 432 | IA32_ENERGY_PERF_BIAS        | Package | See Table 2-2.  |
| 1B1H             | 433 | IA32_PACKAGE_THERM_STATUS    | Package | See Table 2-2.  |
| 1B2H             | 434 | IA32_PACKAGE_THERM_INTERRUPT | Package | See Table 2-2.  |
| 1C8H             | 456 | MSR_LBR_SELECT               | Thread  | Last Branch Record Filtering Select Register (R/W)<br>See Section 17.9.2, "Filtering of Last Branch Records."   |
|                  |     | 0                            |         | CPL_EQ_0  |
|                  |     | 1                            |         | CPL_NEQ_0   |
|                  |     | 2                            |         | JCC   |
|                  |     | 3                            |         | NEAR_REL_CALL   |
|                  |     | 4                            |         | NEAR_IND_CALL   |
|                  |     | 5                            |         | NEAR_RET  |
|                  |     | 6                            |         | NEAR_IND_JMP  |
|                  |     | 7                            |         | NEAR_REL_JMP  |
|                  |     | 8                            |         | FAR_BRANCH  |
| 63:9             |     | Reserved                     |         |   |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS           | Thread  | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-3) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_0_FROM_IP (at 680H).  |
| 1D9H             | 473 | IA32_DEBUGCTL                | Thread  | Debug Control (R/W)<br>See Table 2-2.   |
|                  |     | 0                            |         | LBR: Last Branch Record   |
|                  |     | 1                            |         | BTF   |
|                  |     | 5:2                          |         | Reserved  |
|                  |     | 6                            |         | TR: Branch Trace  |



**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 7                          |        | BTS: Log Branch Trace Message to BTS buffer  |
|                  |     | 8                          |        | BTINT  |
|                  |     | 9                          |        | BTS_OFF_OS   |
|                  |     | 10                         |        | BTS_OFF_USER   |
|                  |     | 11                         |        | FREEZE_LBR_ON_PMI  |
|                  |     | 12                         |        | FREEZE_PERFMON_ON_PMI  |
|                  |     | 13                         |        | ENABLE_UNCORE_PMI  |
|                  |     | 14                         |        | FREEZE_WHILE_SMM   |
|                  |     | 63:15                      |        | Reserved   |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Thread | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.                       |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Thread | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled. |
| 1F2H             | 498 | IA32_SMRR_PHYSBASE         | Core   | See Table 2-2.   |
| 1F3H             | 499 | IA32_SMRR_PHYSMASK         | Core   | See Table 2-2.   |
| 1FCH             | 508 | MSR_POWER_CTL              | Core   | See <a href="http://biosbits.org">http://biosbits.org</a> .  |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Thread | See Table 2-2.   |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Thread | See Table 2-2.   |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Thread | See Table 2-2.   |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Thread | See Table 2-2.   |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Thread | See Table 2-2.   |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Thread | See Table 2-2.   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Thread | See Table 2-2.   |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Thread | See Table 2-2.   |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Thread | See Table 2-2.   |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Thread | See Table 2-2.   |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Thread | See Table 2-2.   |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Thread | See Table 2-2.   |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Thread | See Table 2-2.   |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6        | Thread | See Table 2-2.   |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7        | Thread | See Table 2-2.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7        | Thread  | See Table 2-2.  |
| 210H             | 528 | IA32_MTRR_PHYSBASE8        | Thread  | See Table 2-2.  |
| 211H             | 529 | IA32_MTRR_PHYSMASK8        | Thread  | See Table 2-2.  |
| 212H             | 530 | IA32_MTRR_PHYSBASE9        | Thread  | See Table 2-2.  |
| 213H             | 531 | IA32_MTRR_PHYSMASK9        | Thread  | See Table 2-2.  |
| 250H             | 592 | IA32_MTRR_FIX64K_00000     | Thread  | See Table 2-2.  |
| 258H             | 600 | IA32_MTRR_FIX16K_80000     | Thread  | See Table 2-2.  |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000     | Thread  | See Table 2-2.  |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000      | Thread  | See Table 2-2.  |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000      | Thread  | See Table 2-2.  |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000      | Thread  | See Table 2-2.  |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000      | Thread  | See Table 2-2.  |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000      | Thread  | See Table 2-2.  |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000      | Thread  | See Table 2-2.  |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000      | Thread  | See Table 2-2.  |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000      | Thread  | See Table 2-2.  |
| 277H             | 631 | IA32_PAT                   | Thread  | See Table 2-2.  |
| 280H             | 640 | IA32_MC0_CTL2              | Core    | See Table 2-2.  |
| 281H             | 641 | IA32_MC1_CTL2              | Core    | See Table 2-2.  |
| 282H             | 642 | IA32_MC2_CTL2              | Core    | See Table 2-2.  |
| 283H             | 643 | IA32_MC3_CTL2              | Core    | See Table 2-2.  |
| 284H             | 644 | IA32_MC4_CTL2              | Package | Always 0 (CMCI not supported).  |
| 2FFH             | 767 | IA32_MTRR_DEF_TYPE         | Thread  | Default Memory Types (R/W)<br>See Table 2-2.                          |
| 309H             | 777 | IA32_FIXED_CTR0            | Thread  | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2. |
| 30AH             | 778 | IA32_FIXED_CTR1            | Thread  | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2. |
| 30BH             | 779 | IA32_FIXED_CTR2            | Thread  | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2. |
| 345H             | 837 | IA32_PERF_CAPABILITIES     | Thread  | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."               |
|                  |     | 5:0                        |         | LBR Format<br>See Table 2-2.  |
|                  |     | 6                          |         | PEBS Record Format.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |        | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|--------|----------------------------|--------|---|
| Hex              | Dec    |                            |        |   |
|                  |        | 7                          |        | PEBSSaveArchRegs<br>See Table 2-2.  |
|                  |        | 11:8                       |        | PEBS_REC_FORMAT<br>See Table 2-2.   |
|                  |        | 12                         |        | SMM_FREEZE<br>See Table 2-2.  |
|                  |        | 63:13                      |        | Reserved  |
| 38DH             | 909    | IA32_FIXED_CTR_CTRL        | Thread | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.           |
| 38EH             | 910    | IA32_PERF_GLOBAL_STATUS    |        | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities." |
|                  |        | 0                          | Thread | Ovf_PMC0  |
|                  |        | 1                          | Thread | Ovf_PMC1  |
|                  |        | 2                          | Thread | Ovf_PMC2  |
|                  |        | 3                          | Thread | Ovf_PMC3  |
|                  |        | 4                          | Core   | Ovf_PMC4 (if CPUID.0AH:EAX[15:8] > 4)                                     |
|                  |        | 5                          | Core   | Ovf_PMC5 (if CPUID.0AH:EAX[15:8] > 5)                                     |
|                  |        | 6                          | Core   | Ovf_PMC6 (if CPUID.0AH:EAX[15:8] > 6)                                     |
|                  |        | 7                          | Core   | Ovf_PMC7 (if CPUID.0AH:EAX[15:8] > 7)                                     |
|                  |        | 31:8                       |        | Reserved  |
|                  |        | 32                         | Thread | Ovf_FixedCtr0   |
|                  |        | 33                         | Thread | Ovf_FixedCtr1   |
|                  |        | 34                         | Thread | Ovf_FixedCtr2   |
|                  |        | 60:35                      |        | Reserved  |
|                  |        | 61                         | Thread | Ovf_Uncore  |
|                  |        | 62                         | Thread | Ovf_BufDSSAVE   |
| 63               | Thread | CondChgd                   |        |   |
| 38FH             | 911    | IA32_PERF_GLOBAL_CTRL      | Thread | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities." |
|                  |        | 0                          | Thread | Set 1 to enable PMC0 to count.  |
|                  |        | 1                          | Thread | Set 1 to enable PMC1 to count.  |
|                  |        | 2                          | Thread | Set 1 to enable PMC2 to count.  |
|                  |        | 3                          | Thread | Set 1 to enable PMC3 to count.  |
|                  |        | 4                          | Core   | Set 1 to enable PMC4 to count (if CPUID.0AH:EAX[15:8] > 4).               |
|                  |        | 5                          | Core   | Set 1 to enable PMC5 to count (if CPUID.0AH:EAX[15:8] > 5).               |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
|                  |     | 6                          | Core   | Set 1 to enable PMC6 to count (if CPUID.0AH:EAX[15:8] > 6).               |
|                  |     | 7                          | Core   | Set 1 to enable PMC7 to count (if CPUID.0AH:EAX[15:8] > 7).               |
|                  |     | 31:8                       |        | Reserved  |
|                  |     | 32                         | Thread | Set 1 to enable FixedCtr0 to count.                                       |
|                  |     | 33                         | Thread | Set 1 to enable FixedCtr1 to count.                                       |
|                  |     | 34                         | Thread | Set 1 to enable FixedCtr2 to count.                                       |
|                  |     | 63:35                      |        | Reserved  |
| 390H             | 912 | IA32_PERF_GLOBAL_OVF_CTRL  |        | See Table 2-2. See Section 18.6.2.2, "Global Counter Control Facilities." |
|                  |     | 0                          | Thread | Set 1 to clear Ovf_PMC0.  |
|                  |     | 1                          | Thread | Set 1 to clear Ovf_PMC1.  |
|                  |     | 2                          | Thread | Set 1 to clear Ovf_PMC2.  |
|                  |     | 3                          | Thread | Set 1 to clear Ovf_PMC3.  |
|                  |     | 4                          | Core   | Set 1 to clear Ovf_PMC4 (if CPUID.0AH:EAX[15:8] > 4).                     |
|                  |     | 5                          | Core   | Set 1 to clear Ovf_PMC5 (if CPUID.0AH:EAX[15:8] > 5).                     |
|                  |     | 6                          | Core   | Set 1 to clear Ovf_PMC6 (if CPUID.0AH:EAX[15:8] > 6).                     |
|                  |     | 7                          | Core   | Set 1 to clear Ovf_PMC7 (if CPUID.0AH:EAX[15:8] > 7).                     |
|                  |     | 31:8                       |        | Reserved  |
|                  |     | 32                         | Thread | Set 1 to clear Ovf_FixedCtr0.   |
|                  |     | 33                         | Thread | Set 1 to clear Ovf_FixedCtr1.   |
|                  |     | 34                         | Thread | Set 1 to clear Ovf_FixedCtr2.   |
|                  |     | 60:35                      |        | Reserved  |
|                  |     | 3F1H                       | 1009   | MSR_PEBS_ENABLE   |
| 0                |     |                            |        | Enable PEBS on IA32_PMC0. (R/W)   |
| 1                |     |                            |        | Enable PEBS on IA32_PMC1. (R/W)   |
| 2                |     |                            |        | Enable PEBS on IA32_PMC2. (R/W)   |
| 3                |     |                            |        | Enable PEBS on IA32_PMC3. (R/W)   |
| 31:4             |     |                            |        | Reserved  |
| 32               |     |                            |        | Enable Load Latency on IA32_PMC0. (R/W)                                   |
| 33               |     |                            |        | Enable Load Latency on IA32_PMC1. (R/W)                                   |
|                  |     | 34                         |        | Enable Load Latency on IA32_PMC2. (R/W)                                   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 35                         |         | Enable Load Latency on IA32_PMC3. (R/W)   |
|                  |      | 62:36                      |         | Reserved  |
|                  |      | 63                         |         | Enable Precise Store (R/W)  |
| 3F6H             | 1014 | MSR_PEBS_LD_LAT            | Thread  | See Section 18.3.1.1.2, "Load Latency Performance Monitoring Facility."   |
|                  |      | 15:0                       |         | Minimum threshold latency value of tagged load operation that will be counted. (R/W)  |
|                  |      | 63:36                      |         | Reserved  |
| 3F8H             | 1016 | MSR_PKG_C3_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                           |
|                  |      | 63:0                       |         | Package C3 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C3 states. Count at the same frequency as the TSC.  |
| 3F9H             | 1017 | MSR_PKG_C6_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                           |
|                  |      | 63:0                       |         | Package C6 Residency Counter. (R/O)<br>Value since last reset that this package is in processor-specific C6 states. Count at the same frequency as the TSC. |
| 3FAH             | 1018 | MSR_PKG_C7_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                           |
|                  |      | 63:0                       |         | Package C7 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C7 states. Count at the same frequency as the TSC.  |
| 3FCH             | 1020 | MSR_CORE_C3_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                           |
|                  |      | 63:0                       |         | CORE C3 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C3 states. Count at the same frequency as the TSC.        |
| 3FDH             | 1021 | MSR_CORE_C6_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                           |
|                  |      | 63:0                       |         | CORE C6 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C6 states. Count at the same frequency as the TSC.        |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|------|----------------------------|-------|--|
| Hex              | Dec  |                            |       |  |
| 3FEH             | 1022 | MSR_CORE_C7_RESIDENCY      | Core  | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                    |
|                  |      | 63:0                       |       | CORE C7 Residency Counter (R/O)<br>Value since last reset that this core is in processor-specific C7 states. Count at the same frequency as the TSC. |
| 400H             | 1024 | IA32_MCO_CTL               | Core  | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 401H             | 1025 | IA32_MCO_STATUS            | Core  | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |
| 402H             | 1026 | IA32_MCO_ADDR              | Core  | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 403H             | 1027 | IA32_MCO_MISC              | Core  | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 404H             | 1028 | IA32_MC1_CTL               | Core  | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 405H             | 1029 | IA32_MC1_STATUS            | Core  | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |
| 406H             | 1030 | IA32_MC1_ADDR              | Core  | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 407H             | 1031 | IA32_MC1_MISC              | Core  | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 408H             | 1032 | IA32_MC2_CTL               | Core  | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 409H             | 1033 | IA32_MC2_STATUS            | Core  | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Core  | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 40BH             | 1035 | IA32_MC2_MISC              | Core  | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 40CH             | 1036 | IA32_MC3_CTL               | Core  | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 40DH             | 1037 | IA32_MC3_STATUS            | Core  | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |
| 40EH             | 1038 | IA32_MC3_ADDR              | Core  | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 40FH             | 1039 | IA32_MC3_MISC              | Core  | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 410H             | 1040 | IA32_MC4_CTL               | Core  | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
|                  |      | 0                          |       | PCU Hardware Error (R/W)<br>When set, enables signaling of PCU hardware detected errors.   |
|                  |      | 1                          |       | PCU Controller Error (R/W)<br>When set, enables signaling of PCU controller detected errors.   |
|                  |      | 2                          |       | PCU Firmware Error (R/W)<br>When set, enables signaling of PCU firmware detected errors.   |
|                  |      | 63:2                       |       | Reserved   |
| 411H             | 1041 | IA32_MC4_STATUS            | Core  | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| 480H             | 1152 | IA32_VMX_BASIC             | Thread | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."                     |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL      | Thread | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls." |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL     | Thread | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."     |
| 483H             | 1155 | IA32_VMX_EXIT_CTL          | Thread | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Table 2-2.<br>See Appendix A.4, "VM-Exit Controls."                     |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL         | Thread | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Table 2-2.<br>See Appendix A.5, "VM-Entry Controls."                   |
| 485H             | 1157 | IA32_VMX_MISC              | Thread | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.6, "Miscellaneous Data."                |
| 486H             | 1158 | IA32_VMX_CR0_FIXED0        | Thread | Capability Reporting Register of CR0 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."             |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1        | Thread | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.7, "VMX-Fixed Bits in CR0."             |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0        | Thread | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."             |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1        | Thread | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.<br>See Appendix A.8, "VMX-Fixed Bits in CR4."             |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM         | Thread | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.<br>See Appendix A.9, "VMCS Enumeration."               |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields   | Scope   | Bit Description  |
|------------------|------|------------------------------|---------|--|
| Hex              | Dec  |                              |         |  |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL2      | Thread  | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls."   |
| 48CH             | 1164 | IA32_VMX_EPT_VPID_ENUM       | Thread  | Capability Reporting Register of EPT and VPID (R/O)<br>See Table 2-2   |
| 48DH             | 1165 | IA32_VMX_TRUE_PINBASED_CTL2  | Thread  | Capability Reporting Register of Pin-Based VM-Execution Flex Controls (R/O)<br>See Table 2-2   |
| 48EH             | 1166 | IA32_VMX_TRUE_PROCBASED_CTL2 | Thread  | Capability Reporting Register of Primary Processor-Based VM-Execution Flex Controls (R/O)<br>See Table 2-2   |
| 48FH             | 1167 | IA32_VMX_TRUE_EXIT_CTL2      | Thread  | Capability Reporting Register of VM-Exit Flex Controls (R/O)<br>See Table 2-2  |
| 490H             | 1168 | IA32_VMX_TRUE_ENTRY_CTL2     | Thread  | Capability Reporting Register of VM-Entry Flex Controls (R/O)<br>See Table 2-2   |
| 4C1H             | 1217 | IA32_A_PMC0                  | Thread  | See Table 2-2.   |
| 4C2H             | 1218 | IA32_A_PMC1                  | Thread  | See Table 2-2.   |
| 4C3H             | 1219 | IA32_A_PMC2                  | Thread  | See Table 2-2.   |
| 4C4H             | 1220 | IA32_A_PMC3                  | Thread  | See Table 2-2.   |
| 4C5H             | 1221 | IA32_A_PMC4                  | Core    | See Table 2-2.   |
| 4C6H             | 1222 | IA32_A_PMC5                  | Core    | See Table 2-2.   |
| 4C7H             | 1223 | IA32_A_PMC6                  | Core    | See Table 2-2.   |
| 4C8H             | 1224 | IA32_A_PMC7                  | Core    | See Table 2-2.   |
| 600H             | 1536 | IA32_DS_AREA                 | Thread  | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT          | Package | Unit Multipliers used in RAPL Interfaces (R/O)<br>See Section 14.9.1, "RAPL Interfaces."   |
| 60AH             | 1546 | MSR_PKGC3_IRTL               | Package | Package C3 Interrupt Response Limit (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 9:0                          |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C3 state.                                 |



**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. The following time unit encodings are supported:<br>000b: 1 ns<br>001b: 32 ns<br>010b: 1024 ns<br>011b: 32768 ns<br>100b: 1048576 ns<br>101b: 33554432 ns  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |
|                  |      | 63:16                      |         | Reserved  |
| 60BH             | 1547 | MSR_PKGC6_IRTL             | Package | Package C6 Interrupt Response Limit (R/W)<br>This MSR defines the budget allocated for the package to exit from a C6 to a C0 state, where an interrupt request can be delivered to the core and serviced. Additional core-exit latency may be applicable depending on the actual C-state the core is in.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states. |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C6 state.  |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. The following time unit encodings are supported:<br>000b: 1 ns<br>001b: 32 ns<br>010b: 1024 ns<br>011b: 32768 ns<br>100b: 1048576 ns<br>101b: 33554432 ns  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 63:16                      |         | Reserved   |
| 60DH             | 1549 | MSR_PKG_C2_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |
|                  |      | 63:0                       |         | Package C2 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C2 states. Count at the same frequency as the TSC.   |
| 610H             | 1552 | MSR_PKG_POWER_LIMIT        | Package | PKG RAPL Power Limit Control (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 611H             | 1553 | MSR_PKG_ENERGY_STATUS      | Package | PKG Energy Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."  |
| 614H             | 1556 | MSR_PKG_POWER_INFO         | Package | PKG RAPL Parameters (R/W) See Section 14.9.3, "Package RAPL Domain."   |
| 638H             | 1592 | MSR_PPO_POWER_LIMIT        | Package | PPO RAPL Power Limit Control (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."  |
| 680H             | 1664 | MSR_LASTBRANCH_0_FROM_IP   | Thread  | Last Branch Record 0 From IP (R/W)<br>One of sixteen pairs of last branch record registers on the last branch record stack. This part of the stack contains pointers to the source instruction. See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.9.1 and record format in Section 17.4.8.1.</li> </ul> |
| 681H             | 1665 | MSR_LASTBRANCH_1_FROM_IP   | Thread  | Last Branch Record 1 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 682H             | 1666 | MSR_LASTBRANCH_2_FROM_IP   | Thread  | Last Branch Record 2 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 683H             | 1667 | MSR_LASTBRANCH_3_FROM_IP   | Thread  | Last Branch Record 3 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 684H             | 1668 | MSR_LASTBRANCH_4_FROM_IP   | Thread  | Last Branch Record 4 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 685H             | 1669 | MSR_LASTBRANCH_5_FROM_IP   | Thread  | Last Branch Record 5 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 686H             | 1670 | MSR_LASTBRANCH_6_FROM_IP   | Thread  | Last Branch Record 6 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 687H             | 1671 | MSR_LASTBRANCH_7_FROM_IP   | Thread  | Last Branch Record 7 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 688H             | 1672 | MSR_LASTBRANCH_8_FROM_IP   | Thread  | Last Branch Record 8 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 689H             | 1673 | MSR_LASTBRANCH_9_FROM_IP   | Thread  | Last Branch Record 9 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |

**Table 2-19. MSRs Supported by Intel® Processors  
based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| 68AH             | 1674 | MSR_LASTBRANCH_10_FROM_IP  | Thread | Last Branch Record 10 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68BH             | 1675 | MSR_LASTBRANCH_11_FROM_IP  | Thread | Last Branch Record 11 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68CH             | 1676 | MSR_LASTBRANCH_12_FROM_IP  | Thread | Last Branch Record 12 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68DH             | 1677 | MSR_LASTBRANCH_13_FROM_IP  | Thread | Last Branch Record 13 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68EH             | 1678 | MSR_LASTBRANCH_14_FROM_IP  | Thread | Last Branch Record 14 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 68FH             | 1679 | MSR_LASTBRANCH_15_FROM_IP  | Thread | Last Branch Record 15 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 6C0H             | 1728 | MSR_LASTBRANCH_0_TO_IP     | Thread | Last Branch Record 0 To IP (R/W)<br>One of sixteen pairs of last branch record registers on the last branch record stack. This part of the stack contains pointers to the destination instruction. |
| 6C1H             | 1729 | MSR_LASTBRANCH_1_TO_IP     | Thread | Last Branch Record 1 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C2H             | 1730 | MSR_LASTBRANCH_2_TO_IP     | Thread | Last Branch Record 2 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C3H             | 1731 | MSR_LASTBRANCH_3_TO_IP     | Thread | Last Branch Record 3 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C4H             | 1732 | MSR_LASTBRANCH_4_TO_IP     | Thread | Last Branch Record 4 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C5H             | 1733 | MSR_LASTBRANCH_5_TO_IP     | Thread | Last Branch Record 5 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C6H             | 1734 | MSR_LASTBRANCH_6_TO_IP     | Thread | Last Branch Record 6 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C7H             | 1735 | MSR_LASTBRANCH_7_TO_IP     | Thread | Last Branch Record 7 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C8H             | 1736 | MSR_LASTBRANCH_8_TO_IP     | Thread | Last Branch Record 8 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6C9H             | 1737 | MSR_LASTBRANCH_9_TO_IP     | Thread | Last Branch Record 9 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.   |
| 6CAH             | 1738 | MSR_LASTBRANCH_10_TO_IP    | Thread | Last Branch Record 10 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6CBH             | 1739 | MSR_LASTBRANCH_11_TO_IP    | Thread | Last Branch Record 11 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |

**Table 2-19. MSRs Supported by Intel® Processors based on Intel® microarchitecture code name Sandy Bridge (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 6CCH             | 1740 | MSR_LASTBRANCH_12_TO_IP    | Thread | Last Branch Record 12 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.                                 |
| 6CDH             | 1741 | MSR_LASTBRANCH_13_TO_IP    | Thread | Last Branch Record 13 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.                                 |
| 6CEH             | 1742 | MSR_LASTBRANCH_14_TO_IP    | Thread | Last Branch Record 14 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.                                 |
| 6CFH             | 1743 | MSR_LASTBRANCH_15_TO_IP    | Thread | Last Branch Record 15 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.                                 |
| 6E0H             | 1760 | IA32_TSC_DEADLINE          | Thread | See Table 2-2.  |
| 802H-83FH        |      | X2APIC MSRs                | Thread | See Table 2-2.  |
| C000_0080H       |      | IA32_EFER                  | Thread | Extended Feature Enables<br>See Table 2-2.  |
| C000_0081H       |      | IA32_STAR                  | Thread | System Call Target Address (R/W)<br>See Table 2-2.  |
| C000_0082H       |      | IA32_LSTAR                 | Thread | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.  |
| C000_0084H       |      | IA32_FMASK                 | Thread | System Call Flag Mask (R/W)<br>See Table 2-2.   |
| C000_0100H       |      | IA32_FS_BASE               | Thread | Map of BASE Address of FS (R/W)<br>See Table 2-2.   |
| C000_0101H       |      | IA32_GS_BASE               | Thread | Map of BASE Address of GS (R/W)<br>See Table 2-2.   |
| C000_0102H       |      | IA32_KERNEL_GS_BASE        | Thread | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.   |
| C000_0103H       |      | IA32_TSC_AUX               | Thread | AUXILIARY TSC Signature (R/W)<br>See Table 2-2 and Section 17.17.2, "IA32_TSC_AUX Register and RDTSCP Support." |

### 2.10.1 MSRs In 2nd Generation Intel® Core™ Processor Family (Based on Intel® Microarchitecture Code Name Sandy Bridge)

Table 2-20 and Table 2-21 list model-specific registers (MSRs) that are specific to the 2nd generation Intel® Core™ processor family (based on Intel microarchitecture code name Sandy Bridge). These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2AH; see Table 2-1.

**Table 2-20. MSRs Supported by 2nd Generation Intel® Core™ Processors  
(Intel® microarchitecture code name Sandy Bridge)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 1ADH             | 429  | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |      | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |      | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |
|                  |      | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |      | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |
|                  |      | 63:32                      |         | Reserved   |
| 60CH             | 1548 | MSR_PKGC7_IRTL             | Package | Package C7 Interrupt Response Limit (R/W)<br>This MSR defines the budget allocated for the package to exit from a C7 to a C0 state, where interrupt request can be delivered to the core and serviced. Additional core-exit latency may be applicable depending on the actual C-state the core is in.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states. |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C7 state.   |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. The following time unit encodings are supported:<br>000b: 1 ns<br>001b: 32 ns<br>010b: 1024 ns<br>011b: 32768 ns<br>100b: 1048576 ns<br>101b: 33554432 ns   |
|                  |      | 14:13                      |         | Reserved   |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.  |
|                  |      | 63:16                      |         | Reserved   |

**Table 2-20. MSRs Supported by 2nd Generation Intel® Core™ Processors (Intel® microarchitecture code name Sandy Bridge) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 639H             | 1593 | MSR_PP0_ENERGY_STATUS      | Package | PP0 Energy Status (R/O)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."            |
| 63AH             | 1594 | MSR_PP0_POLICY             | Package | PP0 Balance Policy (R/W)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."           |
| 640H             | 1600 | MSR_PP1_POWER_LIMIT        | Package | PP1 RAPL Power Limit Control (R/W)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains." |
| 641H             | 1601 | MSR_PP1_ENERGY_STATUS      | Package | PP1 Energy Status (R/O)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."            |
| 642H             | 1602 | MSR_PP1_POLICY             | Package | PP1 Balance Policy (R/W)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains."           |

See Table 2-19, Table 2-20, and Table 2-21 for MSR definitions applicable to processors with CPUID signature 06\_2AH.

Table 2-21 lists the MSRs of uncore PMU for Intel processors with CPUID signature 06\_2AH.

**Table 2-21. Uncore PMU MSRs Supported by 2nd Generation Intel® Core™ Processors**

| Register Address |       | Register Name / Bit Fields | Scope    | Bit Description                           |
|------------------|-------|----------------------------|----------|---|
| Hex              | Dec   |                            |          |   |
| 391H             | 913   | MSR_UNC_PERF_GLOBAL_CTRL   | Package  | Uncore PMU Global Control                 |
|                  |       | 0                          |          | Slice 0 select.                           |
|                  |       | 1                          |          | Slice 1 select.                           |
|                  |       | 2                          |          | Slice 2 select.                           |
|                  |       | 3                          |          | Slice 3 select.                           |
|                  |       | 4                          |          | Slice 4 select.                           |
|                  |       | 18:5                       |          | Reserved                                  |
|                  |       | 29                         |          | Enable all uncore counters.               |
|                  |       | 30                         |          | Enable wake on PMI.                       |
|                  |       | 31                         |          | Enable Freezing counter when overflow.    |
|                  | 63:32 |                            | Reserved |   |
| 392H             | 914   | MSR_UNC_PERF_GLOBAL_STATUS | Package  | Uncore PMU Main Status                    |
|                  |       | 0                          |          | Fixed counter overflowed.                 |
|                  |       | 1                          |          | An ARB counter overflowed.                |
|                  |       | 2                          |          | Reserved                                  |
|                  |       | 3                          |          | A CBox counter overflowed (on any slice). |
|                  |       | 63:4                       |          | Reserved                                  |
| 394H             | 916   | MSR_UNC_PERF_FIXED_CTRL    | Package  | Uncore Fixed Counter Control (R/W)        |
|                  |       | 19:0                       |          | Reserved                                  |
|                  |       | 20                         |          | Enable overflow propagation.              |

Table 2-21. Uncore PMU MSRs Supported by 2nd Generation Intel® Core™ Processors

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 21                         |         | Reserved  |
|                  |      | 22                         |         | Enable counting.  |
|                  |      | 63:23                      |         | Reserved  |
| 395H             | 917  | MSR_UNC_PERF_FIXED_CTR     | Package | Uncore Fixed Counter  |
|                  |      | 47:0                       |         | Current count.  |
|                  |      | 63:48                      |         | Reserved  |
| 396H             | 918  | MSR_UNC_CBO_CONFIG         | Package | Uncore C-Box Configuration Information (R/O)  |
|                  |      | 3:0                        |         | Report the number of C-Box units with performance counters, including processor cores and processor graphics. |
|                  |      | 63:4                       |         | Reserved  |
| 3B0H             | 946  | MSR_UNC_ARB_PERFCTR0       | Package | Uncore Arb Unit, Performance Counter 0  |
| 3B1H             | 947  | MSR_UNC_ARB_PERFCTR1       | Package | Uncore Arb Unit, Performance Counter 1  |
| 3B2H             | 944  | MSR_UNC_ARB_PERFEVTSELO    | Package | Uncore Arb Unit, Counter 0 Event Select MSR   |
| 3B3H             | 945  | MSR_UNC_ARB_PERFEVTSEL1    | Package | Uncore Arb unit, Counter 1 Event Select MSR   |
| 700H             | 1792 | MSR_UNC_CBO_0_PERFEVTSELO  | Package | Uncore C-Box 0, Counter 0 Event Select MSR  |
| 701H             | 1793 | MSR_UNC_CBO_0_PERFEVTSEL1  | Package | Uncore C-Box 0, Counter 1 Event Select MSR  |
| 702H             | 1794 | MSR_UNC_CBO_0_PERFEVTSEL2  | Package | Uncore C-Box 0, Counter 2 Event Select MSR  |
| 703H             | 1795 | MSR_UNC_CBO_0_PERFEVTSEL3  | Package | Uncore C-Box 0, Counter 3 Event Select MSR  |
| 705H             | 1797 | MSR_UNC_CBO_0_UNIT_STATUS  | Package | Uncore C-Box 0, Unit Status for Counter 0-3   |
| 706H             | 1798 | MSR_UNC_CBO_0_PERFCTR0     | Package | Uncore C-Box 0, Performance Counter 0   |
| 707H             | 1799 | MSR_UNC_CBO_0_PERFCTR1     | Package | Uncore C-Box 0, Performance Counter 1   |
| 708H             | 1800 | MSR_UNC_CBO_0_PERFCTR2     | Package | Uncore C-Box 0, Performance Counter 2   |
| 709H             | 1801 | MSR_UNC_CBO_0_PERFCTR3     | Package | Uncore C-Box 0, Performance Counter 3   |
| 710H             | 1808 | MSR_UNC_CBO_1_PERFEVTSELO  | Package | Uncore C-Box 1, Counter 0 Event Select MSR  |
| 711H             | 1809 | MSR_UNC_CBO_1_PERFEVTSEL1  | Package | Uncore C-Box 1, Counter 1 Event Select MSR  |
| 712H             | 1810 | MSR_UNC_CBO_1_PERFEVTSEL2  | Package | Uncore C-Box 1, Counter 2 Event Select MSR  |
| 713H             | 1811 | MSR_UNC_CBO_1_PERFEVTSEL3  | Package | Uncore C-Box 1, Counter 3 Event Select MSR  |
| 715H             | 1813 | MSR_UNC_CBO_1_UNIT_STATUS  | Package | Uncore C-Box 1, Unit Status for Counter 0-3   |
| 716H             | 1814 | MSR_UNC_CBO_1_PERFCTR0     | Package | Uncore C-Box 1, Performance Counter 0   |
| 717H             | 1815 | MSR_UNC_CBO_1_PERFCTR1     | Package | Uncore C-Box 1, Performance Counter 1   |
| 718H             | 1816 | MSR_UNC_CBO_1_PERFCTR2     | Package | Uncore C-Box 1, Performance Counter 2   |
| 719H             | 1817 | MSR_UNC_CBO_1_PERFCTR3     | Package | Uncore C-Box 1, Performance Counter 3   |
| 720H             | 1824 | MSR_UNC_CBO_2_PERFEVTSELO  | Package | Uncore C-Box 2, Counter 0 Event Select MSR  |
| 721H             | 1825 | MSR_UNC_CBO_2_PERFEVTSEL1  | Package | Uncore C-Box 2, Counter 1 Event Select MSR  |
| 722H             | 1826 | MSR_UNC_CBO_2_PERFEVTSEL2  | Package | Uncore C-Box 2, Counter 2 Event Select MSR  |

**Table 2-21. Uncore PMU MSRs Supported by 2nd Generation Intel® Core™ Processors**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description                             |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 723H             | 1827 | MSR_UNC_CBO_2_PERFEVTSEL3  | Package | Uncore C-Box 2, Counter 3 Event Select MSR  |
| 725H             | 1829 | MSR_UNC_CBO_2_UNIT_STATUS  | Package | Uncore C-Box 2, Unit Status for Counter 0-3 |
| 726H             | 1830 | MSR_UNC_CBO_2_PERFCTR0     | Package | Uncore C-Box 2, Performance Counter 0       |
| 727H             | 1831 | MSR_UNC_CBO_2_PERFCTR1     | Package | Uncore C-Box 2, Performance Counter 1       |
| 728H             | 1832 | MSR_UNC_CBO_3_PERFCTR2     | Package | Uncore C-Box 3, Performance Counter 2       |
| 729H             | 1833 | MSR_UNC_CBO_3_PERFCTR3     | Package | Uncore C-Box 3, Performance Counter 3       |
| 730H             | 1840 | MSR_UNC_CBO_3_PERFEVTSELO  | Package | Uncore C-Box 3, Counter 0 Event Select MSR  |
| 731H             | 1841 | MSR_UNC_CBO_3_PERFEVTSEL1  | Package | Uncore C-Box 3, Counter 1 Event Select MSR  |
| 732H             | 1842 | MSR_UNC_CBO_3_PERFEVTSEL2  | Package | Uncore C-Box 3, Counter 2 Event Select MSR  |
| 733H             | 1843 | MSR_UNC_CBO_3_PERFEVTSEL3  | Package | Uncore C-Box 3, counter 3 Event Select MSR  |
| 735H             | 1845 | MSR_UNC_CBO_3_UNIT_STATUS  | Package | Uncore C-Box 3, Unit Status for Counter 0-3 |
| 736H             | 1846 | MSR_UNC_CBO_3_PERFCTR0     | Package | Uncore C-Box 3, Performance Counter 0       |
| 737H             | 1847 | MSR_UNC_CBO_3_PERFCTR1     | Package | Uncore C-Box 3, Performance Counter 1       |
| 738H             | 1848 | MSR_UNC_CBO_3_PERFCTR2     | Package | Uncore C-Box 3, Performance Counter 2       |
| 739H             | 1849 | MSR_UNC_CBO_3_PERFCTR3     | Package | Uncore C-Box 3, Performance Counter 3       |
| 740H             | 1856 | MSR_UNC_CBO_4_PERFEVTSELO  | Package | Uncore C-Box 4, Counter 0 Event Select MSR  |
| 741H             | 1857 | MSR_UNC_CBO_4_PERFEVTSEL1  | Package | Uncore C-Box 4, Counter 1 Event Select MSR  |
| 742H             | 1858 | MSR_UNC_CBO_4_PERFEVTSEL2  | Package | Uncore C-Box 4, Counter 2 Event Select MSR  |
| 743H             | 1859 | MSR_UNC_CBO_4_PERFEVTSEL3  | Package | Uncore C-Box 4, Counter 3 Event Select MSR  |
| 745H             | 1861 | MSR_UNC_CBO_4_UNIT_STATUS  | Package | Uncore C-Box 4, Unit status for Counter 0-3 |
| 746H             | 1862 | MSR_UNC_CBO_4_PERFCTR0     | Package | Uncore C-Box 4, Performance Counter 0       |
| 747H             | 1863 | MSR_UNC_CBO_4_PERFCTR1     | Package | Uncore C-Box 4, Performance Counter 1       |
| 748H             | 1864 | MSR_UNC_CBO_4_PERFCTR2     | Package | Uncore C-Box 4, Performance Counter 2       |
| 749H             | 1865 | MSR_UNC_CBO_4_PERFCTR3     | Package | Uncore C-Box 4, Performance Counter 3       |

### 2.10.2 MSRs In Intel® Xeon® Processor E5 Family (Based on Intel® Microarchitecture Code Name Sandy Bridge)

Table 2-22 lists additional model-specific registers (MSRs) that are specific to the Intel® Xeon® Processor E5 Family (based on Intel® microarchitecture code name Sandy Bridge). These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2DH, and also supports MSRs listed in Table 2-19 and Table 2-23.

**Table 2-22. Selected MSRs Supported by Intel® Xeon® Processors E5 Family (based on Sandy Bridge microarchitecture)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                   |
|------------------|-----|----------------------------|---------|-----------------------------------|
| Hex              | Dec |                            |         |                                   |
| 17FH             | 383 | MSR_ERROR_CONTROL          | Package | MC Bank Error Configuration (R/W) |
|                  |     | 0                          |         | Reserved                          |



**Table 2-22. Selected MSRs Supported by Intel® Xeon® Processors E5 Family  
(based on Sandy Bridge microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 1                          |         | MemError Log Enable (R/W)<br>When set, enables IMC status bank to log additional info in bits 36:32.       |
|                  |     | 63:2                       |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0,<br>RW if MSR_PLATFORM_INFO.[28] = 1 |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.                                  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 cores active.                                 |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 cores active.                                 |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 cores active.                                 |
|                  |     | 39:32                      | Package | Maximum Ratio Limit for 5C<br>Maximum turbo ratio limit of 5 cores active.                                 |
|                  |     | 47:40                      | Package | Maximum Ratio Limit for 6C<br>Maximum turbo ratio limit of 6 cores active.                                 |
|                  |     | 55:48                      | Package | Maximum Ratio Limit for 7C<br>Maximum turbo ratio limit of 7 cores active.                                 |
|                  |     | 63:56                      | Package | Maximum Ratio Limit for 8C<br>Maximum turbo ratio limit of 8 cores active.                                 |
| 285H             | 645 | IA32_MC5_CTL2              | Package | See Table 2-2.   |
| 286H             | 646 | IA32_MC6_CTL2              | Package | See Table 2-2.   |
| 287H             | 647 | IA32_MC7_CTL2              | Package | See Table 2-2.   |
| 288H             | 648 | IA32_MC8_CTL2              | Package | See Table 2-2.   |
| 289H             | 649 | IA32_MC9_CTL2              | Package | See Table 2-2.   |
| 28AH             | 650 | IA32_MC10_CTL2             | Package | See Table 2-2.   |
| 28BH             | 651 | IA32_MC11_CTL2             | Package | See Table 2-2.   |
| 28CH             | 652 | IA32_MC12_CTL2             | Package | See Table 2-2.   |
| 28DH             | 653 | IA32_MC13_CTL2             | Package | See Table 2-2.   |
| 28EH             | 654 | IA32_MC14_CTL2             | Package | See Table 2-2.   |
| 28FH             | 655 | IA32_MC15_CTL2             | Package | See Table 2-2.   |
| 290H             | 656 | IA32_MC16_CTL2             | Package | See Table 2-2.   |
| 291H             | 657 | IA32_MC17_CTL2             | Package | See Table 2-2.   |
| 292H             | 658 | IA32_MC18_CTL2             | Package | See Table 2-2.   |
| 293H             | 659 | IA32_MC19_CTL2             | Package | See Table 2-2.   |

**Table 2-22. Selected MSRs Supported by Intel® Xeon® Processors E5 Family (based on Sandy Bridge microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 39CH             | 924  | MSR_PEBS_NUM_ALT           | Package | ENABLE_PEBS_NUM_ALT (RW)   |
|                  |      | 0                          |         | ENABLE_PEBS_NUM_ALT (RW)<br>Write 1 to enable alternate PEBS counting logic for specific events requiring additional configuration, see Table 19-17. |
|                  |      | 63:1                       |         | Reserved, must be zero.  |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 415H             | 1045 | IA32_MC5_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 416H             | 1046 | IA32_MC5_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 417H             | 1047 | IA32_MC5_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 419H             | 1049 | IA32_MC6_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 41BH             | 1051 | IA32_MC6_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 41FH             | 1055 | IA32_MC7_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 421H             | 1057 | IA32_MC8_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 422H             | 1058 | IA32_MC8_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 423H             | 1059 | IA32_MC8_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 425H             | 1061 | IA32_MC9_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 426H             | 1062 | IA32_MC9_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 427H             | 1063 | IA32_MC9_MISC              | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 429H             | 1065 | IA32_MC10_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."  |
| 42BH             | 1067 | IA32_MC10_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."  |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs" and Chapter 16.   |

**Table 2-22. Selected MSRs Supported by Intel® Xeon® Processors E5 Family  
(based on Sandy Bridge microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 42FH             | 1071 | IA32_MC11_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 431H             | 1073 | IA32_MC12_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 432H             | 1074 | IA32_MC12_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 433H             | 1075 | IA32_MC12_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 435H             | 1077 | IA32_MC13_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 436H             | 1078 | IA32_MC13_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 437H             | 1079 | IA32_MC13_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 439H             | 1081 | IA32_MC14_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 43BH             | 1083 | IA32_MC14_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 43FH             | 1087 | IA32_MC15_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 441H             | 1089 | IA32_MC16_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 442H             | 1090 | IA32_MC16_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 443H             | 1091 | IA32_MC16_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 445H             | 1093 | IA32_MC17_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 446H             | 1094 | IA32_MC17_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 447H             | 1095 | IA32_MC17_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |
| 449H             | 1097 | IA32_MC18_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16. |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."                  |
| 44BH             | 1099 | IA32_MC18_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."                  |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."                   |

**Table 2-22. Selected MSRs Supported by Intel® Xeon® Processors E5 Family (based on Sandy Bridge microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS" and Chapter 16.                        |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRS."   |
| 44FH             | 1103 | IA32_MC19_MISC             | Package | See Section 15.3.2.4, "IA32_MCi_MISC MSRS."   |
| 613H             | 1555 | MSR_PKG_PERF_STATUS        | Package | Package RAPL Perf Status (R/O)  |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."      |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."                 |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain." |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."               |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS      | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."              |

See Table 2-19, Table 2-22, and Table 2-23 for MSR definitions applicable to processors with CPUID signature 06\_2DH.

### 2.10.3 Additional Uncore PMU MSRs in the Intel® Xeon® Processor E5 Family

Intel Xeon Processor E5 family is based on the Sandy Bridge microarchitecture. The MSR-based uncore PMU interfaces are listed in Table 2-23. For complete detail of the uncore PMU, refer to Intel Xeon Processor E5 Product Family Uncore Performance Monitoring Guide. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_2DH

**Table 2-23. Uncore PMU MSRs in Intel® Xeon® Processor E5 Family**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                                       |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| C08H             |     | MSR_U_PMON_UCLK_FIXED_CTL  | Package | Uncore U-box UCLK Fixed Counter Control               |
| C09H             |     | MSR_U_PMON_UCLK_FIXED_CTR  | Package | Uncore U-box UCLK Fixed Counter                       |
| C10H             |     | MSR_U_PMON_EVNTSELO        | Package | Uncore U-box Perfmon Event Select for U-box Counter 0 |
| C11H             |     | MSR_U_PMON_EVNTSEL1        | Package | Uncore U-box Perfmon Event Select for U-box Counter 1 |
| C16H             |     | MSR_U_PMON_CTR0            | Package | Uncore U-box Perfmon Counter 0                        |
| C17H             |     | MSR_U_PMON_CTR1            | Package | Uncore U-box Perfmon Counter 1                        |
| C24H             |     | MSR_PCU_PMON_BOX_CTL       | Package | Uncore PCU Perfmon for PCU-box-wide Control           |
| C30H             |     | MSR_PCU_PMON_EVNTSELO      | Package | Uncore PCU Perfmon Event Select for PCU Counter 0     |
| C31H             |     | MSR_PCU_PMON_EVNTSEL1      | Package | Uncore PCU Perfmon Event Select for PCU Counter 1     |
| C32H             |     | MSR_PCU_PMON_EVNTSEL2      | Package | Uncore PCU Perfmon Event Select for PCU Counter 2     |

Table 2-23. Uncore PMU MSRs in Intel® Xeon® Processor E5 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| C33H             |     | MSR_PCU_PMON_EVNTSEL3      | Package | Uncore PCU Perfmon Event Select for PCU Counter 3         |
| C34H             |     | MSR_PCU_PMON_BOX_FILTER    | Package | Uncore PCU Perfmon box-wide Filter                        |
| C36H             |     | MSR_PCU_PMON_CTR0          | Package | Uncore PCU Perfmon Counter 0                              |
| C37H             |     | MSR_PCU_PMON_CTR1          | Package | Uncore PCU Perfmon Counter 1                              |
| C38H             |     | MSR_PCU_PMON_CTR2          | Package | Uncore PCU Perfmon Counter 2                              |
| C39H             |     | MSR_PCU_PMON_CTR3          | Package | Uncore PCU Perfmon Counter 3                              |
| D04H             |     | MSR_C0_PMON_BOX_CTL        | Package | Uncore C-box 0 Perfmon Local Box Wide Control             |
| D10H             |     | MSR_C0_PMON_EVNTSEL0       | Package | Uncore C-box 0 Perfmon Event Select for C-box 0 Counter 0 |
| D11H             |     | MSR_C0_PMON_EVNTSEL1       | Package | Uncore C-box 0 Perfmon Event Select for C-box 0 Counter 1 |
| D12H             |     | MSR_C0_PMON_EVNTSEL2       | Package | Uncore C-box 0 Perfmon Event Select for C-box 0 Counter 2 |
| D13H             |     | MSR_C0_PMON_EVNTSEL3       | Package | Uncore C-box 0 Perfmon Event Select for C-box 0 Counter 3 |
| D14H             |     | MSR_C0_PMON_BOX_FILTER     | Package | Uncore C-box 0 Perfmon Box Wide Filter                    |
| D16H             |     | MSR_C0_PMON_CTR0           | Package | Uncore C-box 0 Perfmon Counter 0                          |
| D17H             |     | MSR_C0_PMON_CTR1           | Package | Uncore C-box 0 Perfmon Counter 1                          |
| D18H             |     | MSR_C0_PMON_CTR2           | Package | Uncore C-box 0 Perfmon Counter 2                          |
| D19H             |     | MSR_C0_PMON_CTR3           | Package | Uncore C-box 0 Perfmon Counter 3                          |
| D24H             |     | MSR_C1_PMON_BOX_CTL        | Package | Uncore C-box 1 Perfmon Local Box Wide Control             |
| D30H             |     | MSR_C1_PMON_EVNTSEL0       | Package | Uncore C-box 1 Perfmon Event Select for C-box 1 Counter 0 |
| D31H             |     | MSR_C1_PMON_EVNTSEL1       | Package | Uncore C-box 1 Perfmon Event Select for C-box 1 Counter 1 |
| D32H             |     | MSR_C1_PMON_EVNTSEL2       | Package | Uncore C-box 1 Perfmon Event Select for C-box 1 Counter 2 |
| D33H             |     | MSR_C1_PMON_EVNTSEL3       | Package | Uncore C-box 1 Perfmon Event Select for C-box 1 Counter 3 |
| D34H             |     | MSR_C1_PMON_BOX_FILTER     | Package | Uncore C-box 1 Perfmon Box Wide Filter                    |
| D36H             |     | MSR_C1_PMON_CTR0           | Package | Uncore C-box 1 Perfmon Counter 0                          |
| D37H             |     | MSR_C1_PMON_CTR1           | Package | Uncore C-box 1 Perfmon Counter 1                          |
| D38H             |     | MSR_C1_PMON_CTR2           | Package | Uncore C-box 1 Perfmon Counter 2                          |
| D39H             |     | MSR_C1_PMON_CTR3           | Package | Uncore C-box 1 Perfmon Counter 3                          |
| D44H             |     | MSR_C2_PMON_BOX_CTL        | Package | Uncore C-box 2 Perfmon Local Box Wide Control             |
| D50H             |     | MSR_C2_PMON_EVNTSEL0       | Package | Uncore C-box 2 Perfmon Event Select for C-box 2 Counter 0 |
| D51H             |     | MSR_C2_PMON_EVNTSEL1       | Package | Uncore C-box 2 Perfmon Event Select for C-box 2 Counter 1 |

Table 2-23. Uncore PMU MSRs in Intel® Xeon® Processor E5 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| D52H             |     | MSR_C2_PMON_EVNTSEL2       | Package | Uncore C-box 2 Perfmon Event Select for C-box 2 Counter 2 |
| D53H             |     | MSR_C2_PMON_EVNTSEL3       | Package | Uncore C-box 2 Perfmon Event Select for C-box 2 Counter 3 |
| D54H             |     | MSR_C2_PMON_BOX_FILTER     | Package | Uncore C-box 2 Perfmon Box Wide Filter                    |
| D56H             |     | MSR_C2_PMON_CTR0           | Package | Uncore C-box 2 Perfmon Counter 0                          |
| D57H             |     | MSR_C2_PMON_CTR1           | Package | Uncore C-box 2 Perfmon Counter 1                          |
| D58H             |     | MSR_C2_PMON_CTR2           | Package | Uncore C-box 2 Perfmon Counter 2                          |
| D59H             |     | MSR_C2_PMON_CTR3           | Package | Uncore C-box 2 Perfmon Counter 3                          |
| D64H             |     | MSR_C3_PMON_BOX_CTL        | Package | Uncore C-box 3 Perfmon Local Box Wide Control             |
| D70H             |     | MSR_C3_PMON_EVNTSELO       | Package | Uncore C-box 3 Perfmon Event Select for C-box 3 Counter 0 |
| D71H             |     | MSR_C3_PMON_EVNTSEL1       | Package | Uncore C-box 3 Perfmon Event Select for C-box 3 Counter 1 |
| D72H             |     | MSR_C3_PMON_EVNTSEL2       | Package | Uncore C-box 3 Perfmon Event Select for C-box 3 Counter 2 |
| D73H             |     | MSR_C3_PMON_EVNTSEL3       | Package | Uncore C-box 3 Perfmon Event Select for C-box 3 Counter 3 |
| D74H             |     | MSR_C3_PMON_BOX_FILTER     | Package | Uncore C-box 3 Perfmon Box Wide Filter                    |
| D76H             |     | MSR_C3_PMON_CTR0           | Package | Uncore C-box 3 Perfmon Counter 0                          |
| D77H             |     | MSR_C3_PMON_CTR1           | Package | Uncore C-box 3 Perfmon Counter 1                          |
| D78H             |     | MSR_C3_PMON_CTR2           | Package | Uncore C-box 3 Perfmon Counter 2                          |
| D79H             |     | MSR_C3_PMON_CTR3           | Package | Uncore C-box 3 Perfmon Counter 3                          |
| D84H             |     | MSR_C4_PMON_BOX_CTL        | Package | Uncore C-box 4 Perfmon Local Box Wide Control             |
| D90H             |     | MSR_C4_PMON_EVNTSELO       | Package | Uncore C-box 4 Perfmon Event Select for C-box 4 Counter 0 |
| D91H             |     | MSR_C4_PMON_EVNTSEL1       | Package | Uncore C-box 4 Perfmon Event Select for C-box 4 Counter 1 |
| D92H             |     | MSR_C4_PMON_EVNTSEL2       | Package | Uncore C-box 4 Perfmon Event Select for C-box 4 Counter 2 |
| D93H             |     | MSR_C4_PMON_EVNTSEL3       | Package | Uncore C-box 4 Perfmon Event Select for C-box 4 Counter 3 |
| D94H             |     | MSR_C4_PMON_BOX_FILTER     | Package | Uncore C-box 4 Perfmon Box Wide Filter                    |
| D96H             |     | MSR_C4_PMON_CTR0           | Package | Uncore C-box 4 Perfmon Counter 0                          |
| D97H             |     | MSR_C4_PMON_CTR1           | Package | Uncore C-box 4 Perfmon Counter 1                          |
| D98H             |     | MSR_C4_PMON_CTR2           | Package | Uncore C-box 4 Perfmon Counter 2                          |
| D99H             |     | MSR_C4_PMON_CTR3           | Package | Uncore C-box 4 Perfmon Counter 3                          |
| DA4H             |     | MSR_C5_PMON_BOX_CTL        | Package | Uncore C-box 5 Perfmon Local Box Wide Control             |
| DB0H             |     | MSR_C5_PMON_EVNTSELO       | Package | Uncore C-box 5 Perfmon Event Select for C-box 5 Counter 0 |

Table 2-23. Uncore PMU MSRs in Intel® Xeon® Processor E5 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| DB1H             |     | MSR_C5_PMON_EVNTSEL1       | Package | Uncore C-box 5 Perfmon Event Select for C-box 5 Counter 1 |
| DB2H             |     | MSR_C5_PMON_EVNTSEL2       | Package | Uncore C-box 5 Perfmon Event Select for C-box 5 Counter 2 |
| DB3H             |     | MSR_C5_PMON_EVNTSEL3       | Package | Uncore C-box 5 Perfmon Event Select for C-box 5 Counter 3 |
| DB4H             |     | MSR_C5_PMON_BOX_FILTER     | Package | Uncore C-box 5 Perfmon Box Wide Filter                    |
| DB6H             |     | MSR_C5_PMON_CTR0           | Package | Uncore C-box 5 Perfmon Counter 0                          |
| DB7H             |     | MSR_C5_PMON_CTR1           | Package | Uncore C-box 5 Perfmon Counter 1                          |
| DB8H             |     | MSR_C5_PMON_CTR2           | Package | Uncore C-box 5 Perfmon Counter 2                          |
| DB9H             |     | MSR_C5_PMON_CTR3           | Package | Uncore C-box 5 Perfmon Counter 3                          |
| DC4H             |     | MSR_C6_PMON_BOX_CTL        | Package | Uncore C-box 6 Perfmon Local Box Wide Control             |
| DD0H             |     | MSR_C6_PMON_EVNTSELO       | Package | Uncore C-box 6 Perfmon Event Select for C-box 6 Counter 0 |
| DD1H             |     | MSR_C6_PMON_EVNTSEL1       | Package | Uncore C-box 6 Perfmon Event Select for C-box 6 Counter 1 |
| DD2H             |     | MSR_C6_PMON_EVNTSEL2       | Package | Uncore C-box 6 Perfmon Event Select for C-box 6 Counter 2 |
| DD3H             |     | MSR_C6_PMON_EVNTSEL3       | Package | Uncore C-box 6 Perfmon Event Select for C-box 6 Counter 3 |
| DD4H             |     | MSR_C6_PMON_BOX_FILTER     | Package | Uncore C-box 6 Perfmon Box Wide Filter                    |
| DD6H             |     | MSR_C6_PMON_CTR0           | Package | Uncore C-box 6 Perfmon Counter 0                          |
| DD7H             |     | MSR_C6_PMON_CTR1           | Package | Uncore C-box 6 Perfmon Counter 1                          |
| DD8H             |     | MSR_C6_PMON_CTR2           | Package | Uncore C-box 6 Perfmon Counter 2                          |
| DD9H             |     | MSR_C6_PMON_CTR3           | Package | Uncore C-box 6 Perfmon Counter 3                          |
| DE4H             |     | MSR_C7_PMON_BOX_CTL        | Package | Uncore C-box 7 Perfmon Local Box Wide Control             |
| DF0H             |     | MSR_C7_PMON_EVNTSELO       | Package | Uncore C-box 7 Perfmon Event Select for C-box 7 Counter 0 |
| DF1H             |     | MSR_C7_PMON_EVNTSEL1       | Package | Uncore C-box 7 Perfmon Event Select for C-box 7 Counter 1 |
| DF2H             |     | MSR_C7_PMON_EVNTSEL2       | Package | Uncore C-box 7 Perfmon Event Select for C-box 7 Counter 2 |
| DF3H             |     | MSR_C7_PMON_EVNTSEL3       | Package | Uncore C-box 7 Perfmon Event Select for C-box 7 Counter 3 |
| DF4H             |     | MSR_C7_PMON_BOX_FILTER     | Package | Uncore C-box 7 Perfmon Box Wide Filter                    |
| DF6H             |     | MSR_C7_PMON_CTR0           | Package | Uncore C-box 7 Perfmon Counter 0                          |
| DF7H             |     | MSR_C7_PMON_CTR1           | Package | Uncore C-box 7 Perfmon Counter 1                          |
| DF8H             |     | MSR_C7_PMON_CTR2           | Package | Uncore C-box 7 Perfmon Counter 2                          |
| DF9H             |     | MSR_C7_PMON_CTR3           | Package | Uncore C-box 7 Perfmon Counter 3                          |

## 2.11 MSRS IN THE 3RD GENERATION INTEL® CORE™ PROCESSOR FAMILY (BASED ON INTEL® MICROARCHITECTURE CODE NAME IVY BRIDGE)

The 3rd generation Intel® Core™ processor family and the Intel® Xeon® processor E3-1200v2 product family (based on Intel microarchitecture code name Ivy Bridge) support the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-21, and Table 2-24. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_3AH.

**Table 2-24. Additional MSRs Supported by 3rd Generation Intel® Core™ Processors (based on Intel® microarchitecture code name Ivy Bridge)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the frequency that invariant TSC runs at. Frequency = ratio * 100 MHz.   |
|                  |     | 27:16                      |         | Reserved   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled. |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates that TDP Limit for Turbo mode is not programmable.               |
|                  |     | 31:30                      |         | Reserved   |
|                  |     | 32                         | Package | Low Power Mode Support (LPM) (R/O)<br>When set to 1, indicates that LPM is supported. When set to 0, indicates LPM is not supported.   |
|                  |     | 34:33                      | Package | Number of ConfigTDP Levels (R/O)<br>00: Only Base TDP level available.<br>01: One additional TDP level available.<br>02: Two additional TDP level available.<br>11: Reserved                                     |
|                  |     | 39:35                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz.   |
|                  |     | 55:48                      | Package | Minimum Operating Ratio (R/O)<br>Contains the minimum supported operating ratio in units of 100 MHz.   |
| 63:56            |     | Reserved                   |         |  |



**Table 2-24. Additional MSRs Supported by 3rd Generation Intel® Core™ Processors  
(based on Intel® microarchitecture code name Ivy Bridge) (Contd.)**

| Register Address |     | Register Name / Bit Fields  | Scope | Bit Description  |
|------------------|-----|---|-------|--|
| Hex              | Dec |   |       |  |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL  | Core  | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0   |       | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-state support)<br>001b: C2<br>010b: C6 no retention<br>011b: C6 retention<br>100b: C7<br>101b: C7s<br>111: No package C-state limit.<br>Note: This field cannot be used to limit package C-state to C3. |
|                  |     | 9:3   |       | Reserved   |
|                  |     | 10  |       | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.  |
|                  |     | 14:11   |       | Reserved   |
|                  |     | 15  |       | CFG Lock (R/WO)<br>When set, locks bits 15:0 of this register until next reset.  |
|                  |     | 24:16   |       | Reserved   |
|                  |     | 25  |       | C3 State Auto Demotion Enable (R/W)<br>When set, the processor will conditionally demote C6/C7 requests to C3 based on uncore auto-demote information.   |
|                  |     | 26  |       | C1 State Auto Demotion Enable (R/W)<br>When set, the processor will conditionally demote C3/C6/C7 requests to C1 based on uncore auto-demote information.  |
|                  |     | 27  |       | Enable C3 Undemotion (R/W)<br>When set, enables undemotion from demoted C3.  |
| 28               |     | Enable C1 Undemotion (R/W)<br>When set, enables undemotion from demoted C1. |       |  |

**Table 2-24. Additional MSRs Supported by 3rd Generation Intel® Core™ Processors (based on Intel® microarchitecture code name Ivy Bridge) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:29                      |         | Reserved  |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS      | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."                                |
| 648H             | 1608 | MSR_CONFIG_TDP_NOMINAL     | Package | Base TDP Ratio (R/O)  |
|                  |      | 7:0                        |         | Config_TDP_Base<br>Base TDP level ratio to be used for this specific processor (in units of 100 MHz). |
|                  |      | 63:8                       |         | Reserved  |
| 649H             | 1609 | MSR_CONFIG_TDP_LEVEL1      | Package | ConfigTDP Level 1 ratio and power level (R/O)   |
|                  |      | 14:0                       |         | PKG_TDP_LVL1<br>Power setting for ConfigTDP Level 1.  |
|                  |      | 15                         |         | Reserved  |
|                  |      | 23:16                      |         | Config_TDP_LVL1_Ratio<br>ConfigTDP level 1 ratio to be used for this specific processor.              |
|                  |      | 31:24                      |         | Reserved  |
|                  |      | 46:32                      |         | PKG_MAX_PWR_LVL1<br>Max Power setting allowed for ConfigTDP Level 1.                                  |
|                  |      | 47                         |         | Reserved  |
|                  |      | 62:48                      |         | PKG_MIN_PWR_LVL1<br>MIN Power setting allowed for ConfigTDP Level 1.                                  |
|                  |      | 63                         |         | Reserved  |
| 64AH             | 1610 | MSR_CONFIG_TDP_LEVEL2      | Package | ConfigTDP Level 2 ratio and power level (R/O)   |
|                  |      | 14:0                       |         | PKG_TDP_LVL2<br>Power setting for ConfigTDP Level 2.  |
|                  |      | 15                         |         | Reserved  |
|                  |      | 23:16                      |         | Config_TDP_LVL2_Ratio<br>ConfigTDP level 2 ratio to be used for this specific processor.              |
|                  |      | 31:24                      |         | Reserved  |
|                  |      | 46:32                      |         | PKG_MAX_PWR_LVL2<br>Max Power setting allowed for ConfigTDP Level 2.                                  |
|                  |      | 47                         |         | Reserved  |
|                  |      | 62:48                      |         | PKG_MIN_PWR_LVL2<br>MIN Power setting allowed for ConfigTDP Level 2.                                  |
|                  |      | 63                         |         | Reserved  |
| 64BH             | 1611 | MSR_CONFIG_TDP_CONTROL     | Package | ConfigTDP Control (R/W)   |

**Table 2-24. Additional MSRs Supported by 3rd Generation Intel® Core™ Processors  
(based on Intel® microarchitecture code name Ivy Bridge) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 1:0                        |         | TDP_LEVEL (RW/L)<br>System BIOS can program this field.   |
|                  |      | 30:2                       |         | Reserved.   |
|                  |      | 31                         |         | Config_TDP_Lock (RW/L)<br>When this bit is set, the content of this register is locked until a reset.             |
|                  |      | 63:32                      |         | Reserved  |
| 64CH             | 1612 | MSR_TURBO_ACTIVATION_RATIO | Package | ConfigTDP Control (R/W)   |
|                  |      | 7:0                        |         | MAX_NON_TURBO_RATIO (RW/L)<br>System BIOS can program this field.   |
|                  |      | 30:8                       |         | Reserved  |
|                  |      | 31                         |         | TURBO_ACTIVATION_RATIO_Lock (RW/L)<br>When this bit is set, the content of this register is locked until a reset. |
|                  |      | 63:32                      |         | Reserved  |

See Table 2-19, Table 2-20 and Table 2-21 for other MSR definitions applicable to processors with CPUID signature 06\_3AH.

### 2.11.1 MSRs In Intel® Xeon® Processor E5 v2 Product Family (Based on Ivy Bridge-E Microarchitecture)

Table 2-25 lists model-specific registers (MSRs) that are specific to the Intel® Xeon® Processor E5 v2 Product Family (based on Ivy Bridge-E microarchitecture). These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_3EH, see Table 2-1. These processors supports the MSR interfaces listed in Table 2-19, and Table 2-25.

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 4EH              | 78  | MSR_PPIN_CTL               | Package | Protected Processor Inventory Number Enable Control (R/W) |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 0                          |         | <p>LockOut (R/WO)</p> <p>Set 1 to prevent further writes to MSR_PPIN_CTL. Writing 1 to MSR_PPIN_CTL[bit 0] is permitted only if MSR_PPIN_CTL[bit 1] is clear. Default is 0.</p> <p>BIOS should provide an opt-in menu to enable the user to turn on MSR_PPIN_CTL[bit 1] for privileged inventory initialization agent to access MSR_PPIN. After reading MSR_PPIN, the privileged inventory initialization agent should write '01b' to MSR_PPIN_CTL to disable further access to MSR_PPIN and prevent unauthorized modification to MSR_PPIN_CTL.</p> |
|                  |     | 1                          |         | <p>Enable_PPIN (R/W)</p> <p>If 1, enables MSR_PPIN to be accessible using RDMSR. Once set, attempt to write 1 to MSR_PPIN_CTL[bit 0] will cause #GP.</p> <p>If 0, an attempt to read MSR_PPIN will cause #GP. Default is 0.</p>   |
|                  |     | 63:2                       |         | Reserved  |
| 4FH              | 79  | MSR_PPIN                   | Package | Protected Processor Inventory Number (R/O)  |
|                  |     | 63:0                       |         | <p>Protected Processor Inventory Number (R/O)</p> <p>A unique value within a given CPUID family/model/stepping signature that a privileged inventory initialization agent can access to identify each physical processor, when access to MSR_PPIN is enabled. Access to MSR_PPIN is permitted only if MSR_PPIN_CTL[bits 1:0] = '10b'.</p>   |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | <p>Platform Information</p> <p>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a>.</p>  |
|                  |     | 7:0                        |         | Reserved  |
|                  |     | 15:8                       | Package | <p>Maximum Non-Turbo Ratio (R/O)</p> <p>This is the ratio of the frequency that invariant TSC runs at. Frequency = ratio * 100 MHz.</p>   |
|                  |     | 22:16                      |         | Reserved  |
|                  |     | 23                         | Package | <p>PPIN_CAP (R/O)</p> <p>When set to 1, indicates that Protected Processor Inventory Number (PPIN) capability can be enabled for a privileged system inventory agent to read PPIN from MSR_PPIN.</p> <p>When set to 0, PPIN capability is not supported. An attempt to access MSR_PPIN_CTL or MSR_PPIN will cause #GP.</p>  |
| 27:24            |     | Reserved                   |         |   |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled.  |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates TDP Limit for Turbo mode is not programmable.   |
|                  |     | 30                         | Package | Programmable TJ OFFSET (R/O)<br>When set to 1, indicates that MSR_TEMPERATURE_TARGET,[27:24] is valid and writable to specify a temperature offset.   |
|                  |     | 39:31                      |         | Reserved  |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz.  |
|                  |     | 63:48                      |         | Reserved  |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-sate support)<br>001b: C2<br>010b: C6 no retention<br>011b: C6 retention<br>100b: C7<br>101b: C7s<br>111: No package C-state limit.<br>Note: This field cannot be used to limit package C-state to C3. |
|                  |     | 9:3                        |         | Reserved  |
|                  |     | 10                         |         | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO register specified by MSR_PMG_IO_CAPTURE_BASE to MWAIT instructions.   |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 14:11                      |         | Reserved  |
|                  |     | 15                         |         | CFG Lock (R/WO)<br>When set, locks bits 15:0 of this register until next reset.   |
|                  |     | 63:16                      |         | Reserved  |
| 179H             | 377 | IA32_MCG_CAP               | Thread  | Global Machine Check Capability (R/O)   |
|                  |     | 7:0                        |         | Count   |
|                  |     | 8                          |         | MCG_CTL_P   |
|                  |     | 9                          |         | MCG_EXT_P   |
|                  |     | 10                         |         | MCP_CMCI_P  |
|                  |     | 11                         |         | MCG_TES_P   |
|                  |     | 15:12                      |         | Reserved  |
|                  |     | 23:16                      |         | MCG_EXT_CNT   |
|                  |     | 24                         |         | MCG_SER_P   |
|                  |     | 25                         |         | Reserved  |
|                  |     | 26                         |         | MCG_ELOG_P  |
|                  |     | 63:27                      |         | Reserved  |
| 17FH             | 383 | MSR_ERROR_CONTROL          | Package | MC Bank Error Configuration (R/W)   |
|                  |     | 0                          |         | Reserved  |
|                  |     | 1                          |         | MemError Log Enable (R/W)<br>When set, enables IMC status bank to log additional info in bits 36:32.  |
|                  |     | 63:2                       |         | Reserved  |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Package | Temperature Target  |
|                  |     | 15:0                       |         | Reserved  |
|                  |     | 23:16                      |         | Temperature Target (RO)<br>The minimum temperature at which PROCHOT# will be asserted. The value is degrees C.  |
|                  |     | 27:24                      |         | TCC Activation Offset (R/W)<br>Specifies a temperature offset in degrees C from the temperature target (bits 23:16). PROCHOT# will assert at the offset target temperature. Write is permitted only if MSR_PLATFORM_INFO.[30] is set. |
|                  |     | 63:28                      |         | Reserved  |
| 1AEH             | 430 | MSR_TURBO_RATIO_LIMIT1     | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.   |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 9C<br>Maximum turbo ratio limit of 9 core active.   |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 15:8                       | Package | Maximum Ratio Limit for 10C<br>Maximum turbo ratio limit of 10 core active.  |
|                  |      | 23:16                      | Package | Maximum Ratio Limit for 11C<br>Maximum turbo ratio limit of 11 core active.  |
|                  |      | 31:24                      | Package | Maximum Ratio Limit for 12C<br>Maximum turbo ratio limit of 12 core active.  |
|                  |      | 63:32                      |         | Reserved   |
| 285H             | 645  | IA32_MC5_CTL2              | Package | See Table 2-2.   |
| 286H             | 646  | IA32_MC6_CTL2              | Package | See Table 2-2.   |
| 287H             | 647  | IA32_MC7_CTL2              | Package | See Table 2-2.   |
| 288H             | 648  | IA32_MC8_CTL2              | Package | See Table 2-2.   |
| 289H             | 649  | IA32_MC9_CTL2              | Package | See Table 2-2.   |
| 28AH             | 650  | IA32_MC10_CTL2             | Package | See Table 2-2.   |
| 28BH             | 651  | IA32_MC11_CTL2             | Package | See Table 2-2.   |
| 28CH             | 652  | IA32_MC12_CTL2             | Package | See Table 2-2.   |
| 28DH             | 653  | IA32_MC13_CTL2             | Package | See Table 2-2.   |
| 28EH             | 654  | IA32_MC14_CTL2             | Package | See Table 2-2.   |
| 28FH             | 655  | IA32_MC15_CTL2             | Package | See Table 2-2.   |
| 290H             | 656  | IA32_MC16_CTL2             | Package | See Table 2-2.   |
| 291H             | 657  | IA32_MC17_CTL2             | Package | See Table 2-2.   |
| 292H             | 658  | IA32_MC18_CTL2             | Package | See Table 2-2.   |
| 293H             | 659  | IA32_MC19_CTL2             | Package | See Table 2-2.   |
| 294H             | 660  | IA32_MC20_CTL2             | Package | See Table 2-2.   |
| 295H             | 661  | IA32_MC21_CTL2             | Package | See Table 2-2.   |
| 296H             | 662  | IA32_MC22_CTL2             | Package | See Table 2-2.   |
| 297H             | 663  | IA32_MC23_CTL2             | Package | See Table 2-2.   |
| 298H             | 664  | IA32_MC24_CTL2             | Package | See Table 2-2.   |
| 299H             | 665  | IA32_MC25_CTL2             | Package | See Table 2-2.   |
| 29AH             | 666  | IA32_MC26_CTL2             | Package | See Table 2-2.   |
| 29BH             | 667  | IA32_MC27_CTL2             | Package | See Table 2-2.   |
| 29CH             | 668  | IA32_MC28_CTL2             | Package | See Table 2-2.   |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs".<br>Bank MC5 reports MC errors from the Intel QPI module. |
| 415H             | 1045 | IA32_MC5_STATUS            | Package |  |
| 416H             | 1046 | IA32_MC5_ADDR              | Package |  |
| 417H             | 1047 | IA32_MC5_MISC              | Package |  |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC6 reports MC errors from the integrated I/O module.                                       |
| 419H             | 1049 | IA32_MC6_STATUS            | Package |   |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package |   |
| 41BH             | 1051 | IA32_MC6_MISC              | Package |   |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC7 and MC 8 report MC errors from the two home agents.                                    |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package |   |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package |   |
| 41FH             | 1055 | IA32_MC7_MISC              | Package |   |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC7 and MC 8 report MC errors from the two home agents.                                    |
| 421H             | 1057 | IA32_MC8_STATUS            | Package |   |
| 422H             | 1058 | IA32_MC8_ADDR              | Package |   |
| 423H             | 1059 | IA32_MC8_MISC              | Package |   |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 425H             | 1061 | IA32_MC9_STATUS            | Package |   |
| 426H             | 1062 | IA32_MC9_ADDR              | Package |   |
| 427H             | 1063 | IA32_MC9_MISC              | Package |   |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 429H             | 1065 | IA32_MC10_STATUS           | Package |   |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package |   |
| 42BH             | 1067 | IA32_MC10_MISC             | Package |   |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."<br>Bank MC11 reports MC errors from a specific channel of the integrated memory controller.  |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package |   |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package |   |
| 42FH             | 1071 | IA32_MC11_MISC             | Package |   |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 431H             | 1073 | IA32_MC12_STATUS           | Package |   |
| 432H             | 1074 | IA32_MC12_ADDR             | Package |   |
| 433H             | 1075 | IA32_MC12_MISC             | Package |   |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 435H             | 1077 | IA32_MC13_STATUS           | Package |   |
| 436H             | 1078 | IA32_MC13_ADDR             | Package |   |
| 437H             | 1079 | IA32_MC13_MISC             | Package |   |



**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.    |
| 439H             | 1081 | IA32_MC14_STATUS           | Package |  |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package |  |
| 43BH             | 1083 | IA32_MC14_MISC             | Package |  |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.    |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package |  |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package |  |
| 43FH             | 1087 | IA32_MC15_MISC             | Package |  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.    |
| 441H             | 1089 | IA32_MC16_STATUS           | Package |  |
| 442H             | 1090 | IA32_MC16_ADDR             | Package |  |
| 443H             | 1091 | IA32_MC16_MISC             | Package |  |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC17 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 445H             | 1093 | IA32_MC17_STATUS           | Package |  |
| 446H             | 1094 | IA32_MC17_ADDR             | Package |  |
| 447H             | 1095 | IA32_MC17_MISC             | Package |  |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC18 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 449H             | 1097 | IA32_MC18_STATUS           | Package |  |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package |  |
| 44BH             | 1099 | IA32_MC18_MISC             | Package |  |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC19 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package |  |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package |  |
| 44FH             | 1103 | IA32_MC19_MISC             | Package |  |
| 450H             | 1104 | IA32_MC20_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."<br>Bank MC20 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3.  |
| 451H             | 1105 | IA32_MC20_STATUS           | Package |  |
| 452H             | 1106 | IA32_MC20_ADDR             | Package |  |
| 453H             | 1107 | IA32_MC20_MISC             | Package |  |
| 454H             | 1108 | IA32_MC21_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC21 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 455H             | 1109 | IA32_MC21_STATUS           | Package |  |
| 456H             | 1110 | IA32_MC21_ADDR             | Package |  |
| 457H             | 1111 | IA32_MC21_MISC             | Package |  |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 458H             | 1112 | IA32_MC22_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC22 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 459H             | 1113 | IA32_MC22_STATUS           | Package |  |
| 45AH             | 1114 | IA32_MC22_ADDR             | Package |  |
| 45BH             | 1115 | IA32_MC22_MISC             | Package |  |
| 45CH             | 1116 | IA32_MC23_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC23 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 45DH             | 1117 | IA32_MC23_STATUS           | Package |  |
| 45EH             | 1118 | IA32_MC23_ADDR             | Package |  |
| 45FH             | 1119 | IA32_MC23_MISC             | Package |  |
| 460H             | 1120 | IA32_MC24_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC24 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 461H             | 1121 | IA32_MC24_STATUS           | Package |  |
| 462H             | 1122 | IA32_MC24_ADDR             | Package |  |
| 463H             | 1123 | IA32_MC24_MISC             | Package |  |
| 464H             | 1124 | IA32_MC25_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC25 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 465H             | 1125 | IA32_MC25_STATUS           | Package |  |
| 466H             | 1126 | IA32_MC25_ADDR             | Package |  |
| 467H             | 1127 | IA32_MC2MISC               | Package |  |
| 468H             | 1128 | IA32_MC26_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC26 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 469H             | 1129 | IA32_MC26_STATUS           | Package |  |
| 46AH             | 1130 | IA32_MC26_ADDR             | Package |  |
| 46BH             | 1131 | IA32_MC26_MISC             | Package |  |
| 46CH             | 1132 | IA32_MC27_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC27 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 46DH             | 1133 | IA32_MC27_STATUS           | Package |  |
| 46EH             | 1134 | IA32_MC27_ADDR             | Package |  |
| 46FH             | 1135 | IA32_MC27_MISC             | Package |  |
| 470H             | 1136 | IA32_MC28_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC28 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 471H             | 1137 | IA32_MC28_STATUS           | Package |  |
| 472H             | 1138 | IA32_MC28_ADDR             | Package |  |
| 473H             | 1139 | IA32_MC28_MISC             | Package |  |
| 613H             | 1555 | MSR_PKG_PERF_STATUS        | Package | Package RAPL Perf Status (R/O)   |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."  |

**Table 2-25. MSRs Supported by Intel® Xeon® Processors E5 v2 Product Family (based on Ivy Bridge-E microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS      | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains." |

See Table 2-19, for other MSR definitions applicable to Intel Xeon processor E5 v2 with CPUID signature 06\_3EH.

### 2.11.2 Additional MSRs Supported by Intel® Xeon® Processor E7 v2 Family

Intel® Xeon® processor E7 v2 family (based on Ivy Bridge-E microarchitecture) with CPUID DisplayFamily\_DisplayModel signature 06\_3EH supports the MSR interfaces listed in Table 2-19, Table 2-25, and Table 2-26.

**Table 2-26. Additional MSRs Supported by Intel® Xeon® Processor E7 v2 Family with DisplayFamily\_DisplayModel Signature 06\_3EH**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2. |
|                  |     | 0                          |        | Lock (R/WL)  |
|                  |     | 1                          |        | Enable VMX Inside SMX Operation (R/WL)                         |
|                  |     | 2                          |        | Enable VMX Outside SMX Operation (R/WL)                        |
|                  |     | 14:8                       |        | SENTER Local Functions Enables (R/WL)                          |
|                  |     | 15                         |        | SENTER Global Functions Enable (R/WL)                          |
|                  |     | 63:16                      |        | Reserved   |
| 179H             | 377 | IA32_MCG_CAP               | Thread | Global Machine Check Capability (R/O)                          |
|                  |     | 7:0                        |        | Count  |
|                  |     | 8                          |        | MCG_CTL_P  |
|                  |     | 9                          |        | MCG_EXT_P  |
|                  |     | 10                         |        | MCP_CMCI_P   |
|                  |     | 11                         |        | MCG_TES_P  |
|                  |     | 15:12                      |        | Reserved   |
|                  |     | 23:16                      |        | MCG_EXT_CNT  |
|                  |     | 24                         |        | MCG_SER_P  |
| 63:25            |     | Reserved                   |        |  |
| 17AH             | 378 | IA32_MCG_STATUS            | Thread | Global Machine Check Status (R/WO)                             |
|                  |     | 0                          |        | RIPV   |
|                  |     | 1                          |        | EIPV   |

**Table 2-26. Additional MSRs Supported by Intel® Xeon® Processor E7 v2 Family with DisplayFamily\_DisplayModel Signature 06\_3EH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 2                          |         | MCIP  |
|                  |      | 3                          |         | LMCE Signaled   |
|                  |      | 63:4                       |         | Reserved  |
| 1AEH             | 430  | MSR_TURBO_RATIO_LIMIT1     | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.   |
|                  |      | 7:0                        | Package | Maximum Ratio Limit for 9C<br>Maximum turbo ratio limit of 9 core active.   |
|                  |      | 15:8                       | Package | Maximum Ratio Limit for 10C<br>Maximum turbo ratio limit of 10core active.  |
|                  |      | 23:16                      | Package | Maximum Ratio Limit for 11C<br>Maximum turbo ratio limit of 11 core active.   |
|                  |      | 31:24                      | Package | Maximum Ratio Limit for 12C<br>Maximum turbo ratio limit of 12 core active.   |
|                  |      | 39:32                      | Package | Maximum Ratio Limit for 13C<br>Maximum turbo ratio limit of 13 core active.   |
|                  |      | 47:40                      | Package | Maximum Ratio Limit for 14C<br>Maximum turbo ratio limit of 14 core active.   |
|                  |      | 55:48                      | Package | Maximum Ratio Limit for 15C<br>Maximum turbo ratio limit of 15 core active.   |
|                  |      | 62:56                      |         | Reserved  |
|                  |      | 63                         | Package | Semaphore for Turbo Ratio Limit Configuration<br>If 1, the processor uses override configuration <sup>1</sup> specified in MSR_TURBO_RATIO_LIMIT and MSR_TURBO_RATIO_LIMIT1.<br>If 0, the processor uses factory-set configuration (Default). |
| 29DH             | 669  | IA32_MC29_CTL2             | Package | See Table 2-2.  |
| 29EH             | 670  | IA32_MC30_CTL2             | Package | See Table 2-2.  |
| 29FH             | 671  | IA32_MC31_CTL2             | Package | See Table 2-2.  |
| 3F1H             | 1009 | MSR_PEBS_ENABLE            | Thread  | See Section 18.3.1.1.1, "Processor Event Based Sampling (PEBS)."  |
|                  |      | 0                          |         | Enable PEBS on IA32_PMC0 (R/W)  |
|                  |      | 1                          |         | Enable PEBS on IA32_PMC1 (R/W)  |
|                  |      | 2                          |         | Enable PEBS on IA32_PMC2 (R/W)  |
|                  |      | 3                          |         | Enable PEBS on IA32_PMC3 (R/W)  |
|                  |      | 31:4                       |         | Reserved  |
|                  |      | 32                         |         | Enable Load Latency on IA32_PMC0 (R/W)  |
|                  |      | 33                         |         | Enable Load Latency on IA32_PMC1 (R/W)  |

**Table 2-26. Additional MSRs Supported by Intel® Xeon® Processor E7 v2 Family with DisplayFamily\_DisplayModel Signature 06\_3EH**

| Register Address   |      | Register Name / Bit Fields | Scope   | Bit Description  |
|--|------|----------------------------|---------|--|
| Hex  | Dec  |                            |         |  |
|  |      | 34                         |         | Enable Load Latency on IA32_PMC2 (R/W)   |
|  |      | 35                         |         | Enable Load Latency on IA32_PMC3 (R/W)   |
|  |      | 63:36                      |         | Reserved   |
| 41BH   | 1051 | IA32_MC6_MISC              | Package | Misc MAC Information of Integrated I/O (R/O)<br>See Section 15.3.2.4.  |
|  |      | 5:0                        |         | Recoverable Address LSB  |
|  |      | 8:6                        |         | Address Mode   |
|  |      | 15:9                       |         | Reserved   |
|  |      | 31:16                      |         | PCI Express Requestor ID   |
|  |      | 39:32                      |         | PCI Express Segment Number   |
|  |      | 63:32                      |         | Reserved   |
| 474H   | 1140 | IA32_MC29_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC29 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 475H   | 1141 | IA32_MC29_STATUS           | Package |  |
| 476H   | 1142 | IA32_MC29_ADDR             | Package |  |
| 477H   | 1143 | IA32_MC29_MISC             | Package |  |
| 478H   | 1144 | IA32_MC30_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC30 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 479H   | 1145 | IA32_MC30_STATUS           | Package |  |
| 47AH   | 1146 | IA32_MC30_ADDR             | Package |  |
| 47BH   | 1147 | IA32_MC30_MISC             | Package |  |
| 47CH   | 1148 | IA32_MC31_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC31 reports MC errors from a specific CBo (core broadcast) and its corresponding slice of L3. |
| 47DH   | 1149 | IA32_MC31_STATUS           | Package |  |
| 47EH   | 1150 | IA32_MC31_ADDR             | Package |  |
| 47FH   | 1147 | IA32_MC31_MISC             | Package |  |
| See Table 2-19, Table 2-25 for other MSR definitions applicable to Intel Xeon processor E7 v2 with CPUID signature 06_3AH. |      |                            |         |  |

**NOTES:**

1. An override configuration lower than the factory-set configuration is always supported. An override configuration higher than the factory-set configuration is dependent on features specific to the processor and the platform.

### 2.11.3 Additional Uncore PMU MSRs in the Intel® Xeon® Processor E5 v2 and E7 v2 Families

Intel Xeon Processor E5 v2 and E7 v2 families are based on the Ivy Bridge-E microarchitecture. The MSR-based uncore PMU interfaces are listed in Table 2-23 and Table 2-27. For complete detail of the uncore PMU, refer to Intel Xeon Processor E5 v2 Product Family Uncore Performance Monitoring Guide. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_3EH.

Table 2-27. Uncore PMU MSRs in Intel® Xeon® Processor E5 v2 and E7 v2 Families

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| C00H             | 3072 | MSR_PMON_GLOBAL_CTL        | Package | Uncore Perfmon Per-Socket Global Control                  |
| C01H             | 3073 | MSR_PMON_GLOBAL_STATUS     | Package | Uncore Perfmon Per-Socket Global Status                   |
| C06H             | 3078 | MSR_PMON_GLOBAL_CONFIG     | Package | Uncore Perfmon Per-Socket Global Configuration            |
| C15H             | 3093 | MSR_U_PMON_BOX_STATUS      | Package | Uncore U-box Perfmon U-Box Wide Status                    |
| C35H             | 3125 | MSR_PCU_PMON_BOX_STATUS    | Package | Uncore PCU Perfmon Box Wide Status                        |
| D1AH             | 3354 | MSR_C0_PMON_BOX_FILTER1    | Package | Uncore C-Box 0 Perfmon Box Wide Filter1                   |
| D3AH             | 3386 | MSR_C1_PMON_BOX_FILTER1    | Package | Uncore C-Box 1 Perfmon Box Wide Filter1                   |
| D5AH             | 3418 | MSR_C2_PMON_BOX_FILTER1    | Package | Uncore C-Box 2 Perfmon Box Wide Filter1                   |
| D7AH             | 3450 | MSR_C3_PMON_BOX_FILTER1    | Package | Uncore C-Box 3 Perfmon Box Wide Filter1                   |
| D9AH             | 3482 | MSR_C4_PMON_BOX_FILTER1    | Package | Uncore C-Box 4 Perfmon Box Wide Filter1                   |
| DBAH             | 3514 | MSR_C5_PMON_BOX_FILTER1    | Package | Uncore C-Box 5 Perfmon Box Wide Filter1                   |
| DDAH             | 3546 | MSR_C6_PMON_BOX_FILTER1    | Package | Uncore C-Box 6 Perfmon Box Wide Filter1                   |
| DFAH             | 3578 | MSR_C7_PMON_BOX_FILTER1    | Package | Uncore C-Box 7 Perfmon Box Wide Filter1                   |
| E04H             | 3588 | MSR_C8_PMON_BOX_CTL        | Package | Uncore C-Box 8 Perfmon Local Box Wide Control             |
| E10H             | 3600 | MSR_C8_PMON_EVNTSEL0       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 0 |
| E11H             | 3601 | MSR_C8_PMON_EVNTSEL1       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 1 |
| E12H             | 3602 | MSR_C8_PMON_EVNTSEL2       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 2 |
| E13H             | 3603 | MSR_C8_PMON_EVNTSEL3       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 3 |
| E14H             | 3604 | MSR_C8_PMON_BOX_FILTER     | Package | Uncore C-Box 8 Perfmon Box Wide Filter                    |
| E16H             | 3606 | MSR_C8_PMON_CTR0           | Package | Uncore C-Box 8 Perfmon Counter 0                          |
| E17H             | 3607 | MSR_C8_PMON_CTR1           | Package | Uncore C-Box 8 Perfmon Counter 1                          |
| E18H             | 3608 | MSR_C8_PMON_CTR2           | Package | Uncore C-Box 8 Perfmon Counter 2                          |
| E19H             | 3609 | MSR_C8_PMON_CTR3           | Package | Uncore C-Box 8 Perfmon Counter 3                          |
| E1AH             | 3610 | MSR_C8_PMON_BOX_FILTER1    | Package | Uncore C-Box 8 Perfmon Box Wide Filter1                   |
| E24H             | 3620 | MSR_C9_PMON_BOX_CTL        | Package | Uncore C-Box 9 Perfmon Local Box Wide Control             |
| E30H             | 3632 | MSR_C9_PMON_EVNTSEL0       | Package | Uncore C-Box 9 Perfmon Event Select for C-box 9 Counter 0 |
| E31H             | 3633 | MSR_C9_PMON_EVNTSEL1       | Package | Uncore C-Box 9 Perfmon Event Select for C-box 9 Counter 1 |
| E32H             | 3634 | MSR_C9_PMON_EVNTSEL2       | Package | Uncore C-Box 9 Perfmon Event Select for C-box 9 Counter 2 |
| E33H             | 3635 | MSR_C9_PMON_EVNTSEL3       | Package | Uncore C-Box 9 Perfmon Event Select for C-box 9 Counter 3 |
| E34H             | 3636 | MSR_C9_PMON_BOX_FILTER     | Package | Uncore C-Box 9 Perfmon Box Wide Filter                    |
| E36H             | 3638 | MSR_C9_PMON_CTR0           | Package | Uncore C-Box 9 Perfmon Counter 0                          |

Table 2-27. Uncore PMU MSRs in Intel® Xeon® Processor E5 v2 and E7 v2 Families (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| E37H             | 3639 | MSR_C9_PMON_CTR1           | Package | Uncore C-Box 9 Perfmon Counter 1                            |
| E38H             | 3640 | MSR_C9_PMON_CTR2           | Package | Uncore C-Box 9 Perfmon Counter 2                            |
| E39H             | 3641 | MSR_C9_PMON_CTR3           | Package | Uncore C-Box 9 Perfmon Counter 3                            |
| E3AH             | 3642 | MSR_C9_PMON_BOX_FILTER1    | Package | Uncore C-Box 9 Perfmon Box Wide Filter1                     |
| E44H             | 3652 | MSR_C10_PMON_BOX_CTL       | Package | Uncore C-Box 10 Perfmon Local Box Wide Control              |
| E50H             | 3664 | MSR_C10_PMON_EVNTSELO      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 0 |
| E51H             | 3665 | MSR_C10_PMON_EVNTSEL1      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 1 |
| E52H             | 3666 | MSR_C10_PMON_EVNTSEL2      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 2 |
| E53H             | 3667 | MSR_C10_PMON_EVNTSEL3      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 3 |
| E54H             | 3668 | MSR_C10_PMON_BOX_FILTER    | Package | Uncore C-Box 10 Perfmon Box Wide Filter                     |
| E56H             | 3670 | MSR_C10_PMON_CTR0          | Package | Uncore C-Box 10 Perfmon Counter 0                           |
| E57H             | 3671 | MSR_C10_PMON_CTR1          | Package | Uncore C-Box 10 Perfmon Counter 1                           |
| E58H             | 3672 | MSR_C10_PMON_CTR2          | Package | Uncore C-Box 10 Perfmon Counter 2                           |
| E59H             | 3673 | MSR_C10_PMON_CTR3          | Package | Uncore C-Box 10 Perfmon Counter 3                           |
| E5AH             | 3674 | MSR_C10_PMON_BOX_FILTER1   | Package | Uncore C-Box 10 Perfmon Box Wide Filter1                    |
| E64H             | 3684 | MSR_C11_PMON_BOX_CTL       | Package | Uncore C-Box 11 Perfmon Local Box Wide Control              |
| E70H             | 3696 | MSR_C11_PMON_EVNTSELO      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 0 |
| E71H             | 3697 | MSR_C11_PMON_EVNTSEL1      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 1 |
| E72H             | 3698 | MSR_C11_PMON_EVNTSEL2      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 2 |
| E73H             | 3699 | MSR_C11_PMON_EVNTSEL3      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 3 |
| E74H             | 3700 | MSR_C11_PMON_BOX_FILTER    | Package | Uncore C-Box 11 Perfmon Box Wide Filter                     |
| E76H             | 3702 | MSR_C11_PMON_CTR0          | Package | Uncore C-Box 11 Perfmon Counter 0                           |
| E77H             | 3703 | MSR_C11_PMON_CTR1          | Package | Uncore C-Box 11 Perfmon Counter 1                           |
| E78H             | 3704 | MSR_C11_PMON_CTR2          | Package | Uncore C-Box 11 Perfmon Counter 2                           |
| E79H             | 3705 | MSR_C11_PMON_CTR3          | Package | Uncore C-Box 11 Perfmon Counter 3                           |
| E7AH             | 3706 | MSR_C11_PMON_BOX_FILTER1   | Package | Uncore C-Box 11 Perfmon Box Wide Filter1                    |
| E84H             | 3716 | MSR_C12_PMON_BOX_CTL       | Package | Uncore C-Box 12 Perfmon Local Box Wide Control              |
| E90H             | 3728 | MSR_C12_PMON_EVNTSELO      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 0 |
| E91H             | 3729 | MSR_C12_PMON_EVNTSEL1      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 1 |

Table 2-27. Uncore PMU MSRs in Intel® Xeon® Processor E5 v2 and E7 v2 Families (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| E92H             | 3730 | MSR_C12_PMON_EVNTSEL2      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 2 |
| E93H             | 3731 | MSR_C12_PMON_EVNTSEL3      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 3 |
| E94H             | 3732 | MSR_C12_PMON_BOX_FILTER    | Package | Uncore C-Box 12 Perfmon Box Wide Filter                     |
| E96H             | 3734 | MSR_C12_PMON_CTR0          | Package | Uncore C-Box 12 Perfmon Counter 0                           |
| E97H             | 3735 | MSR_C12_PMON_CTR1          | Package | Uncore C-Box 12 Perfmon Counter 1                           |
| E98H             | 3736 | MSR_C12_PMON_CTR2          | Package | Uncore C-Box 12 Perfmon Counter 2                           |
| E99H             | 3737 | MSR_C12_PMON_CTR3          | Package | Uncore C-Box 12 Perfmon Counter 3                           |
| E9AH             | 3738 | MSR_C12_PMON_BOX_FILTER1   | Package | Uncore C-Box 12 Perfmon Box Wide Filter1                    |
| EA4H             | 3748 | MSR_C13_PMON_BOX_CTL       | Package | Uncore C-Box 13 Perfmon Local Box Wide Control              |
| EBOH             | 3760 | MSR_C13_PMON_EVNTSELO      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 0 |
| EB1H             | 3761 | MSR_C13_PMON_EVNTSEL1      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 1 |
| EB2H             | 3762 | MSR_C13_PMON_EVNTSEL2      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 2 |
| EB3H             | 3763 | MSR_C13_PMON_EVNTSEL3      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 3 |
| EB4H             | 3764 | MSR_C13_PMON_BOX_FILTER    | Package | Uncore C-Box 13 Perfmon Box Wide Filter                     |
| EB6H             | 3766 | MSR_C13_PMON_CTR0          | Package | Uncore C-Box 13 Perfmon Counter 0                           |
| EB7H             | 3767 | MSR_C13_PMON_CTR1          | Package | Uncore C-Box 13 Perfmon Counter 1                           |
| EB8H             | 3768 | MSR_C13_PMON_CTR2          | Package | Uncore C-Box 13 Perfmon Counter 2                           |
| EB9H             | 3769 | MSR_C13_PMON_CTR3          | Package | Uncore C-Box 13 Perfmon Counter 3                           |
| EBAH             | 3770 | MSR_C13_PMON_BOX_FILTER1   | Package | Uncore C-Box 13 Perfmon Box Wide Filter1                    |
| EC4H             | 3780 | MSR_C14_PMON_BOX_CTL       | Package | Uncore C-Box 14 Perfmon Local Box Wide Control              |
| EDO H            | 3792 | MSR_C14_PMON_EVNTSELO      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 0 |
| ED1H             | 3793 | MSR_C14_PMON_EVNTSEL1      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 1 |
| ED2H             | 3794 | MSR_C14_PMON_EVNTSEL2      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 2 |
| ED3H             | 3795 | MSR_C14_PMON_EVNTSEL3      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 3 |
| ED4H             | 3796 | MSR_C14_PMON_BOX_FILTER    | Package | Uncore C-Box 14 Perfmon Box Wide Filter                     |
| ED6H             | 3798 | MSR_C14_PMON_CTR0          | Package | Uncore C-Box 14 Perfmon Counter 0                           |
| ED7H             | 3799 | MSR_C14_PMON_CTR1          | Package | Uncore C-Box 14 Perfmon Counter 1                           |
| ED8H             | 3800 | MSR_C14_PMON_CTR2          | Package | Uncore C-Box 14 Perfmon Counter 2                           |
| ED9H             | 3801 | MSR_C14_PMON_CTR3          | Package | Uncore C-Box 14 Perfmon Counter 3                           |
| EDA H            | 3802 | MSR_C14_PMON_BOX_FILTER1   | Package | Uncore C-Box 14 Perfmon Box Wide Filter1                    |



## 2.12 MSRS IN THE 4TH GENERATION INTEL® CORE™ PROCESSORS (BASED ON HASWELL MICROARCHITECTURE)

The 4th generation Intel® Core™ processor family and Intel® Xeon® processor E3-1200v3 product family (based on Haswell microarchitecture), with CPUID DisplayFamily\_DisplayModel signature 06\_3CH/06\_45H/06\_46H, support the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-21, and Table 2-28. For an MSR listed in Table 2-19 that also appears in Table 2-28, Table 2-28 supercede Table 2-19.

The MSRs listed in Table 2-28 also apply to processors based on Haswell-E microarchitecture (see Section 2.13).

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |       | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-------|----------------------------|--|--|
| Hex              | Dec   |                            |  |  |
| 3BH              | 59    | IA32_TSC_ADJUST            | Thread   | Per-Logical-Processor TSC ADJUST (R/W)<br>See Table 2-2.   |
| CEH              | 206   | MSR_PLATFORM_INFO          | Package  | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |       | 7:0                        |  | Reserved   |
|                  |       | 15:8                       | Package  | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the frequency that invariant TSC runs at. Frequency = ratio * 100 MHz.   |
|                  |       | 27:16                      |  | Reserved   |
|                  |       | 28                         | Package  | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled. |
|                  |       | 29                         | Package  | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates TDP Limit for Turbo mode is not programmable.                    |
|                  |       | 31:30                      |  | Reserved   |
|                  |       | 32                         | Package  | Low Power Mode Support (LPM) (R/O)<br>When set to 1, indicates that LPM is supported. When set to 0, indicates LPM is not supported.   |
|                  |       | 34:33                      | Package  | Number of ConfigTDP Levels (R/O)<br>00: Only Base TDP level available.<br>01: One additional TDP level available.<br>02: Two additional TDP level available.<br>11: Reserved                                     |
|                  |       | 39:35                      |  | Reserved   |
|                  | 47:40 | Package                    | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz. |  |

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 55:48                      | Package | Minimum Operating Ratio (R/O)<br>Contains the minimum supported operating ratio in units of 100 MHz.  |
|                  |     | 63:56                      |         | Reserved  |
| 186H             | 390 | IA32_PERFEVTSELO           | Thread  | Performance Event Select for Counter 0 (R/W)<br>Supports all fields described in Table 2-2 and the fields below.  |
|                  |     | 32                         |         | IN_TX: See Section 18.3.6.5.1.<br>When IN_TX (bit 32) is set, AnyThread (bit 21) should be cleared to prevent incorrect results.  |
| 187H             | 391 | IA32_PERFEVTSEL1           | Thread  | Performance Event Select for Counter 1 (R/W)<br>Supports all fields described in Table 2-2 and the fields below.  |
|                  |     | 32                         |         | IN_TX: See Section 18.3.6.5.1.<br>When IN_TX (bit 32) is set, AnyThread (bit 21) should be cleared to prevent incorrect results.  |
| 188H             | 392 | IA32_PERFEVTSEL2           | Thread  | Performance Event Select for Counter 2 (R/W)<br>Supports all fields described in Table 2-2 and the fields below.  |
|                  |     | 32                         |         | IN_TX: See Section 18.3.6.5.1.<br>When IN_TX (bit 32) is set, AnyThread (bit 21) should be cleared to prevent incorrect results.  |
|                  |     | 33                         |         | IN_TXCP: See Section 18.3.6.5.1.<br>When IN_TXCP=1 & IN_TX=1 and in sampling, a spurious PMI may occur and transactions may continuously abort near overflow conditions. Software should favor using IN_TXCP for counting over sampling. If sampling, software should use large "sample-after" value after clearing the counter configured to use IN_TXCP and also always reset the counter even when no overflow condition was reported. |
| 189H             | 393 | IA32_PERFEVTSEL3           | Thread  | Performance Event Select for Counter 3 (R/W)<br>Supports all fields described in Table 2-2 and the fields below.  |
|                  |     | 32                         |         | IN_TX: See Section 18.3.6.5.1<br>When IN_TX (bit 32) is set, AnyThread (bit 21) should be cleared to prevent incorrect results.   |
| 1C8H             | 456 | MSR_LBR_SELECT             | Thread  | Last Branch Record Filtering Select Register (R/W)  |
|                  |     | 0                          |         | CPL_EQ_0  |
|                  |     | 1                          |         | CPL_NEQ_0   |
|                  |     | 2                          |         | JCC   |
|                  |     | 3                          |         | NEAR_REL_CALL   |

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 4                          |         | NEAR_IND_CALL  |
|                  |      | 5                          |         | NEAR_RET   |
|                  |      | 6                          |         | NEAR_IND_JMP   |
|                  |      | 7                          |         | NEAR_REL_JMP   |
|                  |      | 8                          |         | FAR_BRANCH   |
|                  |      | 9                          |         | EN_CALL_STACK  |
|                  |      | 63:9                       |         | Reserved   |
| 1D9H             | 473  | IA32_DEBUGCTL              | Thread  | Debug Control (R/W)<br>See Table 2-2.  |
|                  |      | 0                          |         | LBR: Last Branch Record  |
|                  |      | 1                          |         | BTF  |
|                  |      | 5:2                        |         | Reserved   |
|                  |      | 6                          |         | TR: Branch Trace   |
|                  |      | 7                          |         | BTS: Log Branch Trace Message to BTS Buffer  |
|                  |      | 8                          |         | BTINT  |
|                  |      | 9                          |         | BTS_OFF_OS   |
|                  |      | 10                         |         | BTS_OFF_USER   |
|                  |      | 11                         |         | FREEZE_LBR_ON_PMI  |
|                  |      | 12                         |         | FREEZE_PERFMON_ON_PMI  |
|                  |      | 13                         |         | ENABLE_UNCORE_PMI  |
|                  |      | 14                         |         | FREEZE_WHILE_SMM   |
|                  |      | 15                         |         | RTM_DEBUG  |
|                  |      | 63:15                      |         | Reserved   |
| 491H             | 1169 | IA32_VMX_VMFUNC            | Thread  | Capability Reporting Register of VM-Function Controls (R/O)<br>See Table 2-2.  |
| 60BH             | 1548 | MSR_PKG_C7_IRTL1           | Package | Package C6/C7 Interrupt Response Limit 1 (R/W)<br>This MSR defines the interrupt response time limit used by the processor to manage a transition to a package C6 or C7 state. The latency programmed in this register is for the shorter-latency sub C-states used by an MWAIT hint to a C6 or C7 state.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 9:0                        |         | Interrupt Response Time Limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C6 or C7 state.   |

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. See Table 2-19 for supported time unit encodings.  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |
|                  |      | 63:16                      |         | Reserved  |
| 60CH             | 1548 | MSR_PKG_C7_IRT2            | Package | Package C6/C7 Interrupt Response Limit 2 (R/W)<br>This MSR defines the interrupt response time limit used by the processor to manage a transition to a package C6 or C7 state. The latency programmed in this register is for the longer-latency sub C-states used by an MWAIT hint to a C6 or C7 state.<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States. |
|                  |      | 9:0                        |         | Interrupt response time limit (R/W)<br>Specifies the limit that should be used to decide if the package should be put into a package C6 or C7 state.  |
|                  |      | 12:10                      |         | Time Unit (R/W)<br>Specifies the encoding value of time unit of the interrupt response time limit. See Table 2-19 for supported time unit encodings.  |
|                  |      | 14:13                      |         | Reserved  |
|                  |      | 15                         |         | Valid (R/W)<br>Indicates whether the values in bits 12:0 are valid and can be used by the processor for package C-state management.   |
|                  |      | 63:16                      |         | Reserved  |
| 613H             | 1555 | MSR_PKG_PERF_STATUS        | Package | PKG Perf Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 648H             | 1608 | MSR_CONFIG_TDP_NOMINAL     | Package | Base TDP Ratio (R/O)  |
|                  |      | 7:0                        |         | Config_TDP_Base<br>Base TDP level ratio to be used for this specific processor (in units of 100 MHz).   |

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:8                       |         | Reserved  |
| 649H             | 1609 | MSR_CONFIG_TDP_LEVEL1      | Package | ConfigTDP Level 1 Ratio and Power Level (R/O)   |
|                  |      | 14:0                       |         | PKG_TDP_LVL1<br>Power setting for ConfigTDP Level 1.  |
|                  |      | 15                         |         | Reserved  |
|                  |      | 23:16                      |         | Config_TDP_LVL1_Ratio<br>ConfigTDP level 1 ratio to be used for this specific processor.              |
|                  |      | 31:24                      |         | Reserved  |
|                  |      | 46:32                      |         | PKG_MAX_PWR_LVL1<br>Max Power setting allowed for ConfigTDP Level 1.                                  |
|                  |      | 62:47                      |         | PKG_MIN_PWR_LVL1<br>MIN Power setting allowed for ConfigTDP Level 1.                                  |
|                  |      | 63                         |         | Reserved  |
| 64AH             | 1610 | MSR_CONFIG_TDP_LEVEL2      | Package | ConfigTDP Level 2 Ratio and Power Level (R/O)   |
|                  |      | 14:0                       |         | PKG_TDP_LVL2<br>Power setting for ConfigTDP Level 2.  |
|                  |      | 15                         |         | Reserved  |
|                  |      | 23:16                      |         | Config_TDP_LVL2_Ratio<br>ConfigTDP level 2 ratio to be used for this specific processor.              |
|                  |      | 31:24                      |         | Reserved  |
|                  |      | 46:32                      |         | PKG_MAX_PWR_LVL2<br>Max Power setting allowed for ConfigTDP Level 2.                                  |
|                  |      | 62:47                      |         | PKG_MIN_PWR_LVL2<br>MIN Power setting allowed for ConfigTDP Level 2.                                  |
|                  |      | 63                         |         | Reserved  |
| 64BH             | 1611 | MSR_CONFIG_TDP_CONTROL     | Package | ConfigTDP Control (R/W)   |
|                  |      | 1:0                        |         | TDP_LEVEL (RW/L)<br>System BIOS can program this field.   |
|                  |      | 30:2                       |         | Reserved  |
|                  |      | 31                         |         | Config_TDP_Lock (RW/L)<br>When this bit is set, the content of this register is locked until a reset. |
|                  |      | 63:32                      |         | Reserved  |
| 64CH             | 1612 | MSR_TURBO_ACTIVATION_RATIO | Package | ConfigTDP Control (R/W)   |

**Table 2-28. Additional MSRs Supported by Processors based on the Haswell or Haswell-E microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 7:0                        |         | MAX_NON_TURBO_RATIO (R/W/L)<br>System BIOS can program this field.   |
|                  |      | 30:8                       |         | Reserved   |
|                  |      | 31                         |         | TURBO_ACTIVATION_RATIO_Lock (R/W/L)<br>When this bit is set, the content of this register is locked until a reset. |
|                  |      | 63:32                      |         | Reserved   |
| C80H             | 3200 | IA32_DEBUG_INTERFACE       | Package | Silicon Debug Feature Control (R/W)<br>See Table 2-2.  |

### 2.12.1 MSRs in 4th Generation Intel® Core™ Processor Family (based on Haswell Microarchitecture)

Table 2-29 lists model-specific registers (MSRs) that are specific to 4th generation Intel® Core™ processor family and Intel® Xeon® processor E3-1200 v3 product family (based on Haswell microarchitecture). These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_3CH/06\_45H/06\_46H, see Table 2-1.

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture)**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core  | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 3:0                        |       | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>0000b: C0/C1 (no package C-state support)<br>0001b: C2<br>0010b: C3<br>0011b: C6<br>0100b: C7<br>0101b: C7s<br>Package C states C7 are not available to processors with a signature of 06_3CH. |
|                  |     | 9:4                        |       | Reserved  |
|                  |     | 10                         |       | I/O MWAIT Redirection Enable (R/W)  |

Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 14:11                      |         | Reserved   |
|                  |     | 15                         |         | CFG Lock (R/WO)  |
|                  |     | 24:16                      |         | Reserved   |
|                  |     | 25                         |         | C3 State Auto Demotion Enable (R/W)  |
|                  |     | 26                         |         | C1 State Auto Demotion Enable (R/W)  |
|                  |     | 27                         |         | Enable C3 Undemotion (R/W)   |
|                  |     | 28                         |         | Enable C1 Undemotion (R/W)   |
|                  |     | 63:29                      |         | Reserved   |
| 17DH             | 390 | MSR_SMM_MCA_CAP            | THREAD  | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability Enhancement. Accessible only while in SMM.  |
|                  |     | 57:0                       |         | Reserved   |
|                  |     | 58                         |         | SMM_Code_Access_Chk (SMM-RO)<br>If set to 1, indicates that the SMM code access restriction is supported and the MSR_SMM_FEATURE_CONTROL is supported. |
|                  |     | 59                         |         | Long_Flow_Indication (SMM-RO)<br>If set to 1, indicates that the SMM long flow indicator is supported and the MSR_SMM_DELAYED is supported.            |
|                  |     | 63:60                      |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |
|                  |     | 63:32                      |         | Reserved   |
| 391H             | 913 | MSR_UNC_PERF_GLOBAL_CTRL   | Package | Uncore PMU Global Control  |
|                  |     | 0                          |         | Core 0 select.   |
|                  |     | 1                          |         | Core 1 select.   |
|                  |     | 2                          |         | Core 2 select.   |
|                  |     | 3                          |         | Core 3 select.   |
|                  |     | 18:4                       |         | Reserved   |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                                |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 29                         |         | Enable all uncore counters.                    |
|                  |     | 30                         |         | Enable wake on PMI.                            |
|                  |     | 31                         |         | Enable Freezing counter when overflow.         |
|                  |     | 63:32                      |         | Reserved                                       |
| 392H             | 914 | MSR_UNC_PERF_GLOBAL_STATUS | Package | Uncore PMU Main Status                         |
|                  |     | 0                          |         | Fixed counter overflowed.                      |
|                  |     | 1                          |         | An ARB counter overflowed.                     |
|                  |     | 2                          |         | Reserved                                       |
|                  |     | 3                          |         | A CBox counter overflowed (on any slice).      |
|                  |     | 63:4                       |         | Reserved                                       |
| 394H             | 916 | MSR_UNC_PERF_FIXED_CTRL    | Package | Uncore Fixed Counter Control (R/W)             |
|                  |     | 19:0                       |         | Reserved                                       |
|                  |     | 20                         |         | Enable overflow propagation.                   |
|                  |     | 21                         |         | Reserved                                       |
|                  |     | 22                         |         | Enable counting.                               |
|                  |     | 63:23                      |         | Reserved                                       |
| 395H             | 917 | MSR_UNC_PERF_FIXED_CTR     | Package | Uncore Fixed Counter                           |
|                  |     | 47:0                       |         | Current count.                                 |
|                  |     | 63:48                      |         | Reserved                                       |
| 396H             | 918 | MSR_UNC_CBO_CONFIG         | Package | Uncore C-Box Configuration Information (R/O)   |
|                  |     | 3:0                        |         | Encoded number of C-Box, derive value by "-1". |
|                  |     | 63:4                       |         | Reserved                                       |
| 3B0H             | 946 | MSR_UNC_ARB_PERFCTR0       | Package | Uncore Arb Unit, Performance Counter 0         |
| 3B1H             | 947 | MSR_UNC_ARB_PERFCTR1       | Package | Uncore Arb Unit, Performance Counter 1         |
| 3B2H             | 944 | MSR_UNC_ARB_PERFEVTSELO    | Package | Uncore Arb Unit, Counter 0 Event Select MSR    |
| 3B3H             | 945 | MSR_UNC_ARB_PERFEVTSEL1    | Package | Uncore Arb Unit, Counter 1 Event Select MSR    |
| 391H             | 913 | MSR_UNC_PERF_GLOBAL_CTRL   | Package | Uncore PMU Global Control                      |
|                  |     | 0                          |         | Core 0 select.                                 |
|                  |     | 1                          |         | Core 1 select.                                 |
|                  |     | 2                          |         | Core 2 select.                                 |
|                  |     | 3                          |         | Core 3 select.                                 |
|                  |     | 18:4                       |         | Reserved                                       |
|                  |     | 29                         |         | Enable all uncore counters.                    |
|                  |     | 30                         |         | Enable wake on PMI.                            |
|                  |     | 31                         |         | Enable Freezing counter when overflow.         |
|                  |     | 63:32                      |         | Reserved                                       |



Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 395H             | 917  | MSR_UNC_PERF_FIXED_CTR     | Package | Uncore Fixed Counter  |
|                  |      | 47:0                       |         | Current count.  |
|                  |      | 63:48                      |         | Reserved  |
| 3B3H             | 945  | MSR_UNC_ARB_PERFEVTSEL1    | Package | Uncore Arb Unit, Counter 1 Event Select MSR   |
| 4E0H             | 1248 | MSR_SMM_FEATURE_CONTROL    | Package | Enhanced SMM Feature Control (SMM-RW)<br>Reports SMM capability Enhancement. Accessible only while in SMM.  |
|                  |      | 0                          |         | Lock (SMM-RWO)<br>When set to '1' locks this register from further changes.   |
|                  |      | 1                          |         | Reserved  |
|                  |      | 2                          |         | SMM_Code_Chk_En (SMM-RW)<br>This control bit is available only if MSR_SMM_MCA_CAP[58] == 1. When set to '0' (default) none of the logical processors are prevented from executing SMM code outside the ranges defined by the SMRR.<br>When set to '1' any logical processor in the package that attempts to execute SMM code not within the ranges defined by the SMRR will assert an unrecoverable MCE.  |
|                  |      | 63:3                       |         | Reserved  |
| 4E2H             | 1250 | MSR_SMM_DELAYED            | Package | SMM Delayed (SMM-RO)<br>Reports the interruptible state of all logical processors in the package. Available only while in SMM and MSR_SMM_MCA_CAP[LONG_FLOW_INDICATION] == 1.   |
|                  |      | N-1:0                      |         | LOG_PROC_STATE (SMM-RO)<br>Each bit represents a logical processor of its state in a long flow of internal operation which delays servicing an interrupt. The corresponding bit will be set at the start of long events such as: Microcode Update Load, C6, WBINVD, Ratio Change, Throttle.<br>The bit is automatically cleared at the end of each long event. The reset value of this field is 0.<br>Only bit positions below N = CPUID.(EAX=0BH, ECX=PKG_LVL):EBX[15:0] can be updated. |
|                  |      | 63:N                       |         | Reserved  |
| 4E3H             | 1251 | MSR_SMM_BLOCKED            | Package | SMM Blocked (SMM-RO)<br>Reports the blocked state of all logical processors in the package. Available only while in SMM.  |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description   |
|------------------|------|-----------------------------|---------|---|
| Hex              | Dec  |                             |         |   |
|                  |      | N-1:0                       |         | LOG_PROC_STATE (SMM-RO)<br>Each bit represents a logical processor of its blocked state to service an SMI. The corresponding bit will be set if the logical processor is in one of the following states: Wait For SIPI or SENTER Sleep.<br>The reset value of this field is 0FFFH.<br>Only bit positions below N = CPUID.(EAX=0BH, ECX=PKG_LVL):EBX[15:0] can be updated. |
|                  |      | 63:N                        |         | Reserved  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT         | Package | Unit Multipliers Used in RAPL Interfaces (R/O)  |
|                  |      | 3:0                         | Package | Power Units<br>See Section 14.9.1, "RAPL Interfaces."   |
|                  |      | 7:4                         | Package | Reserved  |
|                  |      | 12:8                        | Package | Energy Status Units<br>Energy related information (in Joules) is based on the multiplier, $1/2^{\text{ESU}}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 0EH (or 61 micro-joules).  |
|                  |      | 15:13                       | Package | Reserved  |
|                  |      | 19:16                       | Package | Time Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 63:20                       |         | Reserved  |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS       | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."  |
| 640H             | 1600 | MSR_PP1_POWER_LIMIT         | Package | PP1 RAPL Power Limit Control (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."   |
| 641H             | 1601 | MSR_PP1_ENERGY_STATUS       | Package | PP1 Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."  |
| 642H             | 1602 | MSR_PP1_POLICY              | Package | PP1 Balance Policy (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."   |
| 690H             | 1680 | MSR_CORE_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.)   |
|                  |      | 0                           |         | PROCHOT Status (RO)<br>When set, processor core frequency is reduced below the operating system request due to assertion of external PROCHOT.   |
|                  |      | 1                           |         | Thermal Status (RO)<br>When set, frequency is reduced below the operating system request due to a thermal event.  |
|                  |      | 3:2                         |         | Reserved  |

Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 4                          |       | Graphics Driver Status (R0)<br>When set, frequency is reduced below the operating system request due to Processor Graphics driver override.   |
|                  |     | 5                          |       | Autonomous Utilization-Based Frequency Control Status (R0)<br>When set, frequency is reduced below the operating system request because the processor has detected that utilization is low.                                       |
|                  |     | 6                          |       | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.  |
|                  |     | 7                          |       | Reserved  |
|                  |     | 8                          |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption).                               |
|                  |     | 9                          |       | Core Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to domain-level power limiting.  |
|                  |     | 10                         |       | Package-Level Power Limiting PL1 Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL1.  |
|                  |     | 11                         |       | Package-Level PL2 Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL2.  |
|                  |     | 12                         |       | Max Turbo Limit Status (R0)<br>When set, frequency is reduced below the operating system request due to multi-core turbo limits.  |
|                  |     | 13                         |       | Turbo Transition Attenuation Status (R0)<br>When set, frequency is reduced below the operating system request due to Turbo transition attenuation. This prevents performance degradation due to frequent operating ratio changes. |
|                  |     | 15:14                      |       | Reserved  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 17                         |       | <p>Thermal Log</p> <p>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>   |
|                  |     | 19:18                      |       | Reserved  |
|                  |     | 20                         |       | <p>Graphics Driver Log</p> <p>When set, indicates that the Graphics Driver Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>   |
|                  |     | 21                         |       | <p>Autonomous Utilization-Based Frequency Control Log</p> <p>When set, indicates that the Autonomous Utilization-Based Frequency Control Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p> |
|                  |     | 22                         |       | <p>VR Therm Alert Log</p> <p>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>   |
|                  |     | 23                         |       | Reserved  |
|                  |     | 24                         |       | <p>Electrical Design Point Log</p> <p>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>   |
|                  |     | 25                         |       | <p>Core Power Limiting Log</p> <p>When set, indicates that the Core Power Limiting Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>   |
|                  |     | 26                         |       | <p>Package-Level PL1 Power Limiting Log</p> <p>When set, indicates that the Package Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.</p> <p>This log bit will remain set until cleared by software writing 0.</p>                             |

Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |      | Register Name / Bit Fields      | Scope   | Bit Description  |
|------------------|------|---------------------------------|---------|--|
| Hex              | Dec  |                                 |         |  |
|                  |      | 27                              |         | Package-Level PL2 Power Limiting Log<br>When set, indicates that the Package Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 28                              |         | Max Turbo Limit Log<br>When set, indicates that the Max Turbo Limit Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                   |
|                  |      | 29                              |         | Turbo Transition Attenuation Log<br>When set, indicates that the Turbo Transition Attenuation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.         |
|                  |      | 63:30                           |         | Reserved   |
| 6B0H             | 1712 | MSR_GRAPHICS_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in the Processor Graphics (R/W)<br>(Frequency refers to processor graphics frequency.)   |
|                  |      | 0                               |         | PROCHOT Status (R0)<br>When set, frequency is reduced below the operating system request due to assertion of external PROCHOT.   |
|                  |      | 1                               |         | Thermal Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal event.   |
|                  |      | 3:2                             |         | Reserved   |
|                  |      | 4                               |         | Graphics Driver Status (R0)<br>When set, frequency is reduced below the operating system request due to Processor Graphics driver override.  |
|                  |      | 5                               |         | Autonomous Utilization-Based Frequency Control Status (R0)<br>When set, frequency is reduced below the operating system request because the processor has detected that utilization is low.  |
|                  |      | 6                               |         | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.   |
|                  |      | 7                               |         | Reserved   |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 8                          |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption).  |
|                  |     | 9                          |       | Graphics Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to domain-level power limiting.   |
|                  |     | 10                         |       | Package-Level Power Limiting PL1 Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL1.   |
|                  |     | 11                         |       | Package-Level PL2 Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL2.   |
|                  |     | 15:12                      |       | Reserved   |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 19:18                      |       | Reserved   |
|                  |     | 20                         |       | Graphics Driver Log<br>When set, indicates that the Graphics Driver Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 21                         |       | Autonomous Utilization-Based Frequency Control Log<br>When set, indicates that the Autonomous Utilization-Based Frequency Control Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |     | 22                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |

Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description  |
|------------------|------|-----------------------------|---------|--|
| Hex              | Dec  |                             |         |  |
|                  |      | 23                          |         | Reserved   |
|                  |      | 24                          |         | Electrical Design Point Log<br>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                       |
|                  |      | 25                          |         | Core Power Limiting Log<br>When set, indicates that the Core Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                           |
|                  |      | 26                          |         | Package-Level PL1 Power Limiting Log<br>When set, indicates that the Package Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 27                          |         | Package-Level PL2 Power Limiting Log<br>When set, indicates that the Package Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 28                          |         | Max Turbo Limit Log<br>When set, indicates that the Max Turbo Limit Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                   |
|                  |      | 29                          |         | Turbo Transition Attenuation Log<br>When set, indicates that the Turbo Transition Attenuation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.         |
|                  |      | 63:30                       |         | Reserved   |
| 6B1H             | 1713 | MSR_RING_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in the Ring Interconnect (R/W)<br>(Frequency refers to ring interconnect in the uncore.)   |
|                  |      | 0                           |         | PROCHOT Status (R0)<br>When set, frequency is reduced below the operating system request due to assertion of external PROCHOT.   |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 1                          |       | Thermal Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal event.   |
|                  |     | 5:2                        |       | Reserved   |
|                  |     | 6                          |       | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.   |
|                  |     | 7                          |       | Reserved   |
|                  |     | 8                          |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption).  |
|                  |     | 9                          |       | Reserved   |
|                  |     | 10                         |       | Package-Level Power Limiting PL1 Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL1.   |
|                  |     | 11                         |       | Package-Level PL2 Power Limiting Status (R0)<br>When set, frequency is reduced below the operating system request due to package-level power limiting PL2.   |
|                  |     | 15:12                      |       | Reserved   |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                 |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                 |
|                  |     | 19:18                      |       | Reserved.  |
|                  |     | 20                         |       | Graphics Driver Log<br>When set, indicates that the Graphics Driver Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |



Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 21                         |       | Autonomous Utilization-Based Frequency Control Log<br>When set, indicates that the Autonomous Utilization-Based Frequency Control Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |     | 22                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 23                         |       | Reserved   |
|                  |     | 24                         |       | Electrical Design Point Log<br>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 25                         |       | Core Power Limiting Log<br>When set, indicates that the Core Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 26                         |       | Package-Level PL1 Power Limiting Log<br>When set, indicates that the Package Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                             |
|                  |     | 27                         |       | Package-Level PL2 Power Limiting Log<br>When set, indicates that the Package Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                             |
|                  |     | 28                         |       | Max Turbo Limit Log<br>When set, indicates that the Max Turbo Limit Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |

**Table 2-29. MSRs Supported by 4th Generation Intel® Core™ Processors (Haswell Microarchitecture) (Contd.)**

| Register Address   |      | Register Name / Bit Fields | Scope   | Bit Description  |
|--|------|----------------------------|---------|--|
| Hex  | Dec  |                            |         |  |
|  |      | 29                         |         | Turbo Transition Attenuation Log<br>When set, indicates that the Turbo Transition Attenuation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|  |      | 63:30                      |         | Reserved   |
| 700H   | 1792 | MSR_UNC_CBO_0_PERFEVTSELO  | Package | Uncore C-Box 0, Counter 0 Event Select MSR   |
| 701H   | 1793 | MSR_UNC_CBO_0_PERFEVTSEL1  | Package | Uncore C-Box 0, Counter 1 Event Select MSR   |
| 706H   | 1798 | MSR_UNC_CBO_0_PERFCTRO     | Package | Uncore C-Box 0, Performance Counter 0  |
| 707H   | 1799 | MSR_UNC_CBO_0_PERFCTR1     | Package | Uncore C-Box 0, Performance Counter 1  |
| 710H   | 1808 | MSR_UNC_CBO_1_PERFEVTSELO  | Package | Uncore C-Box 1, Counter 0 Event Select MSR   |
| 711H   | 1809 | MSR_UNC_CBO_1_PERFEVTSEL1  | Package | Uncore C-Box 1, Counter 1 Event Select MSR   |
| 716H   | 1814 | MSR_UNC_CBO_1_PERFCTRO     | Package | Uncore C-Box 1, Performance Counter 0  |
| 717H   | 1815 | MSR_UNC_CBO_1_PERFCTR1     | Package | Uncore C-Box 1, Performance Counter 1  |
| 720H   | 1824 | MSR_UNC_CBO_2_PERFEVTSELO  | Package | Uncore C-Box 2, Counter 0 Event Select MSR   |
| 721H   | 1824 | MSR_UNC_CBO_2_PERFEVTSEL1  | Package | Uncore C-Box 2, Counter 1 Event Select MSR   |
| 726H   | 1830 | MSR_UNC_CBO_2_PERFCTRO     | Package | Uncore C-Box 2, Performance Counter 0  |
| 727H   | 1831 | MSR_UNC_CBO_2_PERFCTR1     | Package | Uncore C-Box 2, Performance Counter 1  |
| 730H   | 1840 | MSR_UNC_CBO_3_PERFEVTSELO  | Package | Uncore C-Box 3, Counter 0 Event Select MSR   |
| 731H   | 1841 | MSR_UNC_CBO_3_PERFEVTSEL1  | Package | Uncore C-Box 3, Counter 1 Event Select MSR   |
| 736H   | 1846 | MSR_UNC_CBO_3_PERFCTRO     | Package | Uncore C-Box 3, Performance Counter 0  |
| 737H   | 1847 | MSR_UNC_CBO_3_PERFCTR1     | Package | Uncore C-Box 3, Performance Counter 1  |
| See Table 2-19, Table 2-20, Table 2-21, Table 2-24, Table 2-28 for other MSR definitions applicable to processors with CPUID signatures 063CH, 06_46H. |      |                            |         |  |

### 2.12.2 Additional Residency MSRs Supported in 4th Generation Intel® Core™ Processors

The 4th generation Intel® Core™ processor family (based on Haswell microarchitecture) with CPUID DisplayFamily\_DisplayModel signature 06\_45H supports the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-28, Table 2-29, and Table 2-30.

**Table 2-30. Additional Residency MSRs Supported by 4th Generation Intel® Core™ Processors with DisplayFamily\_DisplayModel Signature 06\_45H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| E2H              | 226  | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |      | 3:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br><br>The following C-state code name encodings are supported:<br>0000b: C0/C1 (no package C-state support)<br>0001b: C2<br>0010b: C3<br>0011b: C6<br>0100b: C7<br>0101b: C7s<br>0110b: C8<br>0111b: C9<br>1000b: C10 |
|                  |      | 9:4                        |         | Reserved   |
|                  |      | 10                         |         | I/O MWAIT Redirection Enable (R/W)   |
|                  |      | 14:11                      |         | Reserved   |
|                  |      | 15                         |         | CFG Lock (R/WO)  |
|                  |      | 24:16                      |         | Reserved   |
|                  |      | 25                         |         | C3 State Auto Demotion Enable (R/W)  |
|                  |      | 26                         |         | C1 State Auto Demotion Enable (R/W)  |
|                  |      | 27                         |         | Enable C3 Undemotion (R/W)   |
|                  |      | 28                         |         | Enable C1 Undemotion (R/W)   |
|                  |      | 63:29                      |         | Reserved   |
| 630H             | 1584 | MSR_PKG_C8_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.  |
|                  |      | 59:0                       |         | Package C8 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C8 states. Count at the same frequency as the TSC.   |
|                  |      | 63:60                      |         | Reserved   |

**Table 2-30. Additional Residency MSRs Supported by 4th Generation Intel® Core™ Processors with DisplayFamily\_DisplayModel Signature 06\_45H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 631H             | 1585 | MSR_PKG_C9_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                            |
|                  |      | 59:0                       |         | Package C9 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C9 states. Count at the same frequency as the TSC.   |
|                  |      | 63:60                      |         | Reserved   |
| 632H             | 1586 | MSR_PKG_C10_RESIDENCY      | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-States.                            |
|                  |      | 59:0                       |         | Package C10 Residency Counter (R/O)<br>Value since last reset that this package is in processor-specific C10 states. Count at the same frequency as the TSC. |
|                  |      | 63:60                      |         | Reserved   |

See Table 2-19, Table 2-20, Table 2-21, Table 2-28, Table 2-29 for other MSR definitions applicable to processors with CPUID signature 06\_45H.

## 2.13 MSRS IN INTEL® XEON® PROCESSOR E5 V3 AND E7 V3 PRODUCT FAMILY

Intel® Xeon® processor E5 v3 family and Intel® Xeon® processor E7 v3 family are based on Haswell-E microarchitecture (CPUID DisplayFamily\_DisplayModel = 06\_3F). These processors supports the MSR interfaces listed in Table 2-19, Table 2-28, and Table 2-31.

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 35H              | 53  | MSR_CORE_THREAD_COUNT      | Package | Configured State of Enabled Processor Core Count and Logical Processor Count (RO) <ul style="list-style-type: none"> <li>After a Power-On RESET, enumerates factory configuration of the number of processor cores and logical processors in the physical package.</li> <li>Following the sequence of (i) BIOS modified a Configuration Mask which selects a subset of processor cores to be active post RESET and (ii) a RESET event after the modification, enumerates the current configuration of enabled processor core count and logical processor count in the physical package.</li> </ul> |
|                  |     | 15:0                       |         | Core_COUNT (RO)<br>The number of processor cores that are currently enabled (by either factory configuration or BIOS configuration) in the physical package.   |

Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 31:16                      |        | THREAD_COUNT (RO)<br>The number of logical processors that are currently enabled (by either factory configuration or BIOS configuration) in the physical package.  |
|                  |     | 63:32                      |        | Reserved   |
| 53H              | 83  | MSR_THREAD_ID_INFO         | Thread | A Hardware Assigned ID for the Logical Processor (RO)  |
|                  |     | 7:0                        |        | Logical_Processor_ID (RO)<br>An implementation-specific numerical value physically assigned to each logical processor. This ID is not related to Initial APIC ID or x2APIC ID, it is unique within a physical package.   |
|                  |     | 63:8                       |        | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core   | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0                        |        | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-state support)<br>001b: C2<br>010b: C6 (non-retention)<br>011b: C6 (retention)<br>111b: No Package C state limits. All C states supported by the processor are available. |
|                  |     | 9:3                        |        | Reserved   |
|                  |     | 10                         |        | I/O MWAIT Redirection Enable (R/W)   |
|                  |     | 14:11                      |        | Reserved   |
|                  |     | 15                         |        | CFG Lock (R/W0)  |
|                  |     | 24:16                      |        | Reserved   |
|                  |     | 25                         |        | C3 State Auto Demotion Enable (R/W)  |
|                  |     | 26                         |        | C1 State Auto Demotion Enable (R/W)  |
|                  |     | 27                         |        | Enable C3 Undemotion (R/W)   |
|                  |     | 28                         |        | Enable C1 Undemotion (R/W)   |
|                  |     | 29                         |        | Package C State Demotion Enable (R/W)  |
|                  |     | 30                         |        | Package C State UnDemotion Enable (R/W)  |
| 63:31            |     | Reserved                   |        |  |

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |       | Register Name / Bit Fields | Scope    | Bit Description   |
|------------------|-------|----------------------------|----------|---|
| Hex              | Dec   |                            |          |   |
| 179H             | 377   | IA32_MCG_CAP               | Thread   | Global Machine Check Capability (R/O)   |
|                  |       | 7:0                        |          | Count   |
|                  |       | 8                          |          | MCG_CTL_P   |
|                  |       | 9                          |          | MCG_EXT_P   |
|                  |       | 10                         |          | MCP_CMCI_P  |
|                  |       | 11                         |          | MCG_TES_P   |
|                  |       | 15:12                      |          | Reserved  |
|                  |       | 23:16                      |          | MCG_EXT_CNT   |
|                  |       | 24                         |          | MCG_SER_P   |
|                  |       | 25                         |          | MCG_EM_P  |
|                  |       | 26                         |          | MCG_ELOG_P  |
|                  | 63:27 |                            | Reserved |   |
| 17DH             | 390   | MSR_SMM_MCA_CAP            | Thread   | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability Enhancement. Accessible only while in SMM.   |
|                  |       | 57:0                       |          | Reserved  |
|                  |       | 58                         |          | SMM_Code_Access_Chk (SMM-RO)<br>If set to 1, indicates that the SMM code access restriction is supported and a host-space interface available to SMM handler. |
|                  |       | 59                         |          | Long_Flow_Indication (SMM-RO)<br>If set to 1, indicates that the SMM long flow indicator is supported and a host-space interface available to SMM handler.    |
|                  |       | 63:60                      |          | Reserved  |
| 17FH             | 383   | MSR_ERROR_CONTROL          | Package  | MC Bank Error Configuration (R/W)   |
|                  |       | 0                          |          | Reserved  |
|                  |       | 1                          |          | MemError Log Enable (R/W)<br>When set, enables IMC status bank to log additional info in bits 36:32.  |
|                  |       | 63:2                       |          | Reserved  |
| 1ADH             | 429   | MSR_TURBO_RATIO_LIMIT      | Package  | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.   |
|                  |       | 7:0                        | Package  | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.   |
|                  |       | 15:8                       | Package  | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.   |
|                  |       | 23:16                      | Package  | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.   |

Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.                                   |
|                  |     | 39:32                      | Package | Maximum Ratio Limit for 5C<br>Maximum turbo ratio limit of 5 core active.                                   |
|                  |     | 47:40                      | Package | Maximum Ratio Limit for 6C<br>Maximum turbo ratio limit of 6 core active.                                   |
|                  |     | 55:48                      | Package | Maximum Ratio Limit for 7C<br>Maximum turbo ratio limit of 7 core active.                                   |
|                  |     | 63:56                      | Package | Maximum Ratio Limit for 8C<br>Maximum turbo ratio limit of 8 core active.                                   |
| 1AEH             | 430 | MSR_TURBO_RATIO_LIMIT1     | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1. |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 9C<br>Maximum turbo ratio limit of 9 core active.                                   |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 10C<br>Maximum turbo ratio limit of 10 core active.                                 |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 11C<br>Maximum turbo ratio limit of 11 core active.                                 |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 12C<br>Maximum turbo ratio limit of 12 core active.                                 |
|                  |     | 39:32                      | Package | Maximum Ratio Limit for 13C<br>Maximum turbo ratio limit of 13 core active.                                 |
|                  |     | 47:40                      | Package | Maximum Ratio Limit for 14C<br>Maximum turbo ratio limit of 14 core active.                                 |
|                  |     | 55:48                      | Package | Maximum Ratio Limit for 15C<br>Maximum turbo ratio limit of 15 core active.                                 |
|                  |     | 63:56                      | Package | Maximum Ratio Limit for 16C<br>Maximum turbo ratio limit of 16 core active.                                 |
| 1AFH             | 431 | MSR_TURBO_RATIO_LIMIT2     | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1. |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 17C<br>Maximum turbo ratio limit of 17 core active.                                 |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 18C<br>Maximum turbo ratio limit of 18 core active.                                 |
|                  |     | 62:16                      | Package | Reserved  |

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63                         | Package | Semaphore for Turbo Ratio Limit Configuration<br>If 1, the processor uses override configuration <sup>1</sup> specified in MSR_TURBO_RATIO_LIMIT, MSR_TURBO_RATIO_LIMIT1 and MSR_TURBO_RATIO_LIMIT2.<br>If 0, the processor uses factory-set configuration (Default). |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC5 reports MC errors from the Intel QPI 0 module.  |
| 415H             | 1045 | IA32_MC5_STATUS            | Package |   |
| 416H             | 1046 | IA32_MC5_ADDR              | Package |   |
| 417H             | 1047 | IA32_MC5_MISC              | Package |   |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC6 reports MC errors from the integrated I/O module.   |
| 419H             | 1049 | IA32_MC6_STATUS            | Package |   |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package |   |
| 41BH             | 1051 | IA32_MC6_MISC              | Package |   |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC7 reports MC errors from the home agent HA 0.   |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package |   |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package |   |
| 41FH             | 1055 | IA32_MC7_MISC              | Package |   |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC8 reports MC errors from the home agent HA 1.   |
| 421H             | 1057 | IA32_MC8_STATUS            | Package |   |
| 422H             | 1058 | IA32_MC8_ADDR              | Package |   |
| 423H             | 1059 | IA32_MC8_MISC              | Package |   |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.   |
| 425H             | 1061 | IA32_MC9_STATUS            | Package |   |
| 426H             | 1062 | IA32_MC9_ADDR              | Package |   |
| 427H             | 1063 | IA32_MC9_MISC              | Package |   |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.   |
| 429H             | 1065 | IA32_MC10_STATUS           | Package |   |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package |   |
| 42BH             | 1067 | IA32_MC10_MISC             | Package |   |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.   |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package |   |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package |   |
| 42FH             | 1071 | IA32_MC11_MISC             | Package |   |



Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 431H             | 1073 | IA32_MC12_STATUS           | Package |  |
| 432H             | 1074 | IA32_MC12_ADDR             | Package |  |
| 433H             | 1075 | IA32_MC12_MISC             | Package |  |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 435H             | 1077 | IA32_MC13_STATUS           | Package |  |
| 436H             | 1078 | IA32_MC13_ADDR             | Package |  |
| 437H             | 1079 | IA32_MC13_MISC             | Package |  |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 439H             | 1081 | IA32_MC14_STATUS           | Package |  |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package |  |
| 43BH             | 1083 | IA32_MC14_MISC             | Package |  |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package |  |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package |  |
| 43FH             | 1087 | IA32_MC15_MISC             | Package |  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 441H             | 1089 | IA32_MC16_STATUS           | Package |  |
| 442H             | 1090 | IA32_MC16_ADDR             | Package |  |
| 443H             | 1091 | IA32_MC16_MISC             | Package |  |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC17 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo0, CBo3, CBo6, CBo9, CBo12, CBo15.  |
| 445H             | 1093 | IA32_MC17_STATUS           | Package |  |
| 446H             | 1094 | IA32_MC17_ADDR             | Package |  |
| 447H             | 1095 | IA32_MC17_MISC             | Package |  |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC18 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo1, CBo4, CBo7, CBo10, CBo13, CBo16. |
| 449H             | 1097 | IA32_MC18_STATUS           | Package |  |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package |  |
| 44BH             | 1099 | IA32_MC18_MISC             | Package |  |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC19 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo2, CBo5, CBo8, CBo11, CBo14, CBo17. |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package |  |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package |  |
| 44FH             | 1103 | IA32_MC19_MISC             | Package |  |
| 450H             | 1104 | IA32_MC20_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC20 reports MC errors from the Intel QPI 1 module.  |
| 451H             | 1105 | IA32_MC20_STATUS           | Package |  |
| 452H             | 1106 | IA32_MC20_ADDR             | Package |  |
| 453H             | 1107 | IA32_MC20_MISC             | Package |  |

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 454H             | 1108 | IA32_MC21_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC21 reports MC errors from the Intel QPI 2 module.  |
| 455H             | 1109 | IA32_MC21_STATUS           | Package |  |
| 456H             | 1110 | IA32_MC21_ADDR             | Package |  |
| 457H             | 1111 | IA32_MC21_MISC             | Package |  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers Used in RAPL Interfaces (R/O)   |
|                  |      | 3:0                        | Package | Power Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 7:4                        | Package | Reserved   |
|                  |      | 12:8                       | Package | Energy Status Units<br>Energy related information (in Joules) is based on the multiplier, $1/2^{ESU}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 0EH (or 61 micro-joules).      |
|                  |      | 15:13                      | Package | Reserved   |
|                  |      | 19:16                      | Package | Time Units<br>See Section 14.9.1, "RAPL Interfaces."   |
|                  |      | 63:20                      |         | Reserved   |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>Energy Consumed by DRAM devices.   |
|                  |      | 31:0                       |         | Energy in 15.3 micro-joules. Requires BIOS configuration to enable DRAM RAPL mode 0 (Direct VR).   |
|                  |      | 63:32                      |         | Reserved   |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 61EH             | 1566 | MSR_PCIE_PLL_RATIO         | Package | Configuration of PCIe PLL Relative to BCLK(R/W)  |
|                  |      | 1:0                        | Package | PCIe Ratio (R/W)<br>00b: Use 5:5 mapping for 100MHz operation (default).<br>01b: Use 5:4 mapping for 125MHz operation.<br>10b: Use 5:3 mapping for 166MHz operation.<br>11b: Use 5:2 mapping for 250MHz operation. |
|                  |      | 2                          | Package | LPLL Select (R/W)<br>If 1, use configured setting of PCIe Ratio.   |
|                  |      | 3                          | Package | LONG RESET (R/W)<br>If 1, wait an additional time-out before re-locking Gen2/Gen3 PLLs.  |
|                  |      | 63:4                       |         | Reserved   |

Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description   |
|------------------|------|-----------------------------|---------|---|
| Hex              | Dec  |                             |         |   |
| 620H             | 1568 | MSR_UNCORE_RATIO_LIMIT      | Package | Uncore Ratio Limit (R/W)<br>Out of reset, the min_ratio and max_ratio fields represent the widest possible range of uncore frequencies. Writing to these fields allows software to control the minimum and the maximum frequency that hardware will select. |
|                  |      | 63:15                       |         | Reserved  |
|                  |      | 14:8                        |         | MIN_RATIO<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.   |
|                  |      | 7                           |         | Reserved  |
|                  |      | 6:0                         |         | MAX_RATIO<br>This field is used to limit the max ratio of the LLC/Ring.   |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS       | Package | Reserved (R/O)<br>Reads return 0.   |
| 690H             | 1680 | MSR_CORE_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.)   |
|                  |      | 0                           |         | PROCHOT Status (R0)<br>When set, processor core frequency is reduced below the operating system request due to assertion of external PROCHOT.   |
|                  |      | 1                           |         | Thermal Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal event.  |
|                  |      | 2                           |         | Power Budget Management Status (R0)<br>When set, frequency is reduced below the operating system request due to PBM limit   |
|                  |      | 3                           |         | Platform Configuration Services Status (R0)<br>When set, frequency is reduced below the operating system request due to PCS limit   |
|                  |      | 4                           |         | Reserved  |
|                  |      | 5                           |         | Autonomous Utilization-Based Frequency Control Status (R0)<br>When set, frequency is reduced below the operating system request because the processor has detected that utilization is low.   |
|                  |      | 6                           |         | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.  |
|                  |      | 7                           |         | Reserved  |

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 8                          |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption).      |
|                  |     | 9                          |       | Reserved   |
|                  |     | 10                         |       | Multi-Core Turbo Status (R0)<br>When set, frequency is reduced below the operating system request due to Multi-Core Turbo limits.  |
|                  |     | 12:11                      |       | Reserved   |
|                  |     | 13                         |       | Core Frequency P1 Status (R0)<br>When set, frequency is reduced below max non-turbo P1.  |
|                  |     | 14                         |       | Core Max N-Core Turbo Frequency Limiting Status (R0)<br>When set, frequency is reduced below max n-core turbo frequency.   |
|                  |     | 15                         |       | Core Frequency Limiting Status (R0)<br>When set, frequency is reduced below the operating system request.  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                     |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                     |
|                  |     | 18                         |       | Power Budget Management Log<br>When set, indicates that the PBM Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.         |
|                  |     | 19                         |       | Platform Configuration Services Log<br>When set, indicates that the PCS Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |     | 20                         |       | Reserved   |

Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
|                  |      | 21                         |        | Autonomous Utilization-Based Frequency Control Log<br>When set, indicates that the AUBFC Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                              |
|                  |      | 22                         |        | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 23                         |        | Reserved   |
|                  |      | 24                         |        | Electrical Design Point Log<br>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 25                         |        | Reserved   |
|                  |      | 26                         |        | Multi-Core Turbo Log<br>When set, indicates that the Multi-Core Turbo Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 28:27                      |        | Reserved   |
|                  |      | 29                         |        | Core Frequency P1 Log<br>When set, indicates that the Core Frequency P1 Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 30                         |        | Core Max N-Core Turbo Frequency Limiting Log<br>When set, indicates that the Core Max n-core Turbo Frequency Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 31                         |        | Core Frequency Limiting Log<br>When set, indicates that the Core Frequency Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                   |
|                  |      | 63:32                      |        | Reserved   |
| C8DH             | 3213 | IA32_QM_EVTSEL             | THREAD | Monitoring Event Select Register (R/W)<br>If CPUID.(EAX=07H, ECX=0):EBX.RDT-M[bit 12] = 1.   |

**Table 2-31. Additional MSRs Supported by Intel® Xeon® Processor E5 v3 Family**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
|                  |      | 7:0                        |        | EventID (Rw)<br>Event encoding:<br>0x0: No monitoring.<br>0x1: L3 occupancy monitoring.<br>All other encoding reserved. |
|                  |      | 31:8                       |        | Reserved  |
|                  |      | 41:32                      |        | RMID (Rw)   |
|                  |      | 63:42                      |        | Reserved  |
| C8EH             | 3214 | IA32_QM_CTR                | THREAD | Monitoring Counter Register (R/O)<br>If CPUID.(EAX=07H, ECX=0):EBX.RDT-M[bit 12] = 1.                                   |
|                  |      | 61:0                       |        | Resource Monitored Data   |
|                  |      | 62                         |        | Unavailable: If 1, indicates data for this RMID is not available or not monitored for this resource or RMID.            |
|                  |      | 63                         |        | Error: If 1, indicates and unsupported RMID or event type was written to IA32_PQR_QM_EVTSEL.                            |
| C8FH             | 3215 | IA32_PQR_ASSOC             | THREAD | Resource Association Register (R/W)   |
|                  |      | 9:0                        |        | RMID  |
|                  |      | 63: 10                     |        | Reserved  |

See Table 2-19, Table 2-28 for other MSR definitions applicable to processors with CPUID signature 06\_3FH.

**NOTES:**

1. An override configuration lower than the factory-set configuration is always supported. An override configuration higher than the factory-set configuration is dependent on features specific to the processor and the platform.

**2.13.1 Additional Uncore PMU MSRs in the Intel® Xeon® Processor E5 v3 Family**

Intel Xeon Processor E5 v3 and E7 v3 family are based on the Haswell-E microarchitecture. The MSR-based uncore PMU interfaces are listed in Table 2-32. For complete detail of the uncore PMU, refer to Intel Xeon Processor E5 v3 Product Family Uncore Performance Monitoring Guide. These processors have a CPUID signature with DisplayFamily\_DisplayModel of 06\_3FH.

**Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                                       |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 700H             |     | MSR_PMON_GLOBAL_CTL        | Package | Uncore Perfmon Per-Socket Global Control              |
| 701H             |     | MSR_PMON_GLOBAL_STATUS     | Package | Uncore Perfmon Per-Socket Global Status               |
| 702H             |     | MSR_PMON_GLOBAL_CONFIG     | Package | Uncore Perfmon Per-Socket Global Configuration        |
| 703H             |     | MSR_U_PMON_UCLK_FIXED_CTL  | Package | Uncore U-Box UCLK Fixed Counter Control               |
| 704H             |     | MSR_U_PMON_UCLK_FIXED_CTR  | Package | Uncore U-Box UCLK Fixed Counter                       |
| 705H             |     | MSR_U_PMON_EVNTSELO        | Package | Uncore U-Box Perfmon Event Select for U-Box Counter 0 |

Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                                       |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 706H             |     | MSR_U_PMON_EVNTSEL1        | Package | Uncore U-Box Perfmon Event Select for U-Box Counter 1 |
| 708H             |     | MSR_U_PMON_BOX_STATUS      | Package | Uncore U-Box Perfmon U-Box Wide Status                |
| 709H             |     | MSR_U_PMON_CTR0            | Package | Uncore U-Box Perfmon Counter 0                        |
| 70AH             |     | MSR_U_PMON_CTR1            | Package | Uncore U-Box Perfmon Counter 1                        |
| 710H             |     | MSR_PCU_PMON_BOX_CTL       | Package | Uncore PCU Perfmon for PCU-Box-Wide Control           |
| 711H             |     | MSR_PCU_PMON_EVNTSELO      | Package | Uncore PCU Perfmon Event Select for PCU Counter 0     |
| 712H             |     | MSR_PCU_PMON_EVNTSEL1      | Package | Uncore PCU Perfmon Event Select for PCU Counter 1     |
| 713H             |     | MSR_PCU_PMON_EVNTSEL2      | Package | Uncore PCU Perfmon Event Select for PCU Counter 2     |
| 714H             |     | MSR_PCU_PMON_EVNTSEL3      | Package | Uncore PCU Perfmon Event Select for PCU Counter 3     |
| 715H             |     | MSR_PCU_PMON_BOX_FILTER    | Package | Uncore PCU Perfmon Box-Wide Filter                    |
| 716H             |     | MSR_PCU_PMON_BOX_STATUS    | Package | Uncore PCU Perfmon Box Wide Status                    |
| 717H             |     | MSR_PCU_PMON_CTR0          | Package | Uncore PCU Perfmon Counter 0                          |
| 718H             |     | MSR_PCU_PMON_CTR1          | Package | Uncore PCU Perfmon Counter 1                          |
| 719H             |     | MSR_PCU_PMON_CTR2          | Package | Uncore PCU Perfmon Counter 2                          |
| 71AH             |     | MSR_PCU_PMON_CTR3          | Package | Uncore PCU Perfmon Counter 3                          |
| 720H             |     | MSR_S0_PMON_BOX_CTL        | Package | Uncore SBo 0 Perfmon for SBo 0 Box-Wide Control       |
| 721H             |     | MSR_S0_PMON_EVNTSELO       | Package | Uncore SBo 0 Perfmon Event Select for SBo 0 Counter 0 |
| 722H             |     | MSR_S0_PMON_EVNTSEL1       | Package | Uncore SBo 0 Perfmon Event Select for SBo 0 Counter 1 |
| 723H             |     | MSR_S0_PMON_EVNTSEL2       | Package | Uncore SBo 0 Perfmon Event Select for SBo 0 Counter 2 |
| 724H             |     | MSR_S0_PMON_EVNTSEL3       | Package | Uncore SBo 0 Perfmon Event Select for SBo 0 Counter 3 |
| 725H             |     | MSR_S0_PMON_BOX_FILTER     | Package | Uncore SBo 0 Perfmon Box-Wide Filter                  |
| 726H             |     | MSR_S0_PMON_CTR0           | Package | Uncore SBo 0 Perfmon Counter 0                        |
| 727H             |     | MSR_S0_PMON_CTR1           | Package | Uncore SBo 0 Perfmon Counter 1                        |
| 728H             |     | MSR_S0_PMON_CTR2           | Package | Uncore SBo 0 Perfmon Counter 2                        |
| 729H             |     | MSR_S0_PMON_CTR3           | Package | Uncore SBo 0 Perfmon Counter 3                        |
| 72AH             |     | MSR_S1_PMON_BOX_CTL        | Package | Uncore SBo 1 Perfmon for SBo 1 Box-Wide Control       |
| 72BH             |     | MSR_S1_PMON_EVNTSELO       | Package | Uncore SBo 1 Perfmon Event Select for SBo 1 Counter 0 |
| 72CH             |     | MSR_S1_PMON_EVNTSEL1       | Package | Uncore SBo 1 Perfmon Event Select for SBo 1 Counter 1 |
| 72DH             |     | MSR_S1_PMON_EVNTSEL2       | Package | Uncore SBo 1 Perfmon Event Select for SBo 1 Counter 2 |
| 72EH             |     | MSR_S1_PMON_EVNTSEL3       | Package | Uncore SBo 1 Perfmon Event Select for SBo 1 Counter 3 |
| 72FH             |     | MSR_S1_PMON_BOX_FILTER     | Package | Uncore SBo 1 Perfmon Box-Wide Filter                  |

**Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 730H             |     | MSR_S1_PMON_CTR0           | Package | Uncore SBo 1 Perfmon Counter 0                            |
| 731H             |     | MSR_S1_PMON_CTR1           | Package | Uncore SBo 1 Perfmon Counter 1                            |
| 732H             |     | MSR_S1_PMON_CTR2           | Package | Uncore SBo 1 Perfmon Counter 2                            |
| 733H             |     | MSR_S1_PMON_CTR3           | Package | Uncore SBo 1 Perfmon Counter 3                            |
| 734H             |     | MSR_S2_PMON_BOX_CTL        | Package | Uncore SBo 2 Perfmon for SBo 2 Box-Wide Control           |
| 735H             |     | MSR_S2_PMON_EVNTSELO       | Package | Uncore SBo 2 Perfmon Event Select for SBo 2 Counter 0     |
| 736H             |     | MSR_S2_PMON_EVNTSEL1       | Package | Uncore SBo 2 Perfmon Event Select for SBo 2 Counter 1     |
| 737H             |     | MSR_S2_PMON_EVNTSEL2       | Package | Uncore SBo 2 Perfmon Event Select for SBo 2 Counter 2     |
| 738H             |     | MSR_S2_PMON_EVNTSEL3       | Package | Uncore SBo 2 Perfmon Event Select for SBo 2 Counter 3     |
| 739H             |     | MSR_S2_PMON_BOX_FILTER     | Package | Uncore SBo 2 Perfmon Box-Wide Filter                      |
| 73AH             |     | MSR_S2_PMON_CTR0           | Package | Uncore SBo 2 Perfmon Counter 0                            |
| 73BH             |     | MSR_S2_PMON_CTR1           | Package | Uncore SBo 2 Perfmon Counter 1                            |
| 73CH             |     | MSR_S2_PMON_CTR2           | Package | Uncore SBo 2 Perfmon Counter 2                            |
| 73DH             |     | MSR_S2_PMON_CTR3           | Package | Uncore SBo 2 Perfmon Counter 3                            |
| 73EH             |     | MSR_S3_PMON_BOX_CTL        | Package | Uncore SBo 3 Perfmon for SBo 3 Box-Wide Control           |
| 73FH             |     | MSR_S3_PMON_EVNTSELO       | Package | Uncore SBo 3 Perfmon Event Select for SBo 3 Counter 0     |
| 740H             |     | MSR_S3_PMON_EVNTSEL1       | Package | Uncore SBo 3 Perfmon Event Select for SBo 3 Counter 1     |
| 741H             |     | MSR_S3_PMON_EVNTSEL2       | Package | Uncore SBo 3 Perfmon Event Select for SBo 3 Counter 2     |
| 742H             |     | MSR_S3_PMON_EVNTSEL3       | Package | Uncore SBo 3 Perfmon Event Select for SBo 3 Counter 3     |
| 743H             |     | MSR_S3_PMON_BOX_FILTER     | Package | Uncore SBo 3 Perfmon Box-Wide Filter                      |
| 744H             |     | MSR_S3_PMON_CTR0           | Package | Uncore SBo 3 Perfmon Counter 0                            |
| 745H             |     | MSR_S3_PMON_CTR1           | Package | Uncore SBo 3 Perfmon Counter 1                            |
| 746H             |     | MSR_S3_PMON_CTR2           | Package | Uncore SBo 3 Perfmon Counter 2                            |
| 747H             |     | MSR_S3_PMON_CTR3           | Package | Uncore SBo 3 Perfmon Counter 3                            |
| E00H             |     | MSR_CO_PMON_BOX_CTL        | Package | Uncore C-Box 0 Perfmon for Box-Wide Control               |
| E01H             |     | MSR_CO_PMON_EVNTSELO       | Package | Uncore C-Box 0 Perfmon Event Select for C-Box 0 Counter 0 |
| E02H             |     | MSR_CO_PMON_EVNTSEL1       | Package | Uncore C-Box 0 Perfmon Event Select for C-Box 0 Counter 1 |
| E03H             |     | MSR_CO_PMON_EVNTSEL2       | Package | Uncore C-Box 0 Perfmon Event Select for C-Box 0 Counter 2 |
| E04H             |     | MSR_CO_PMON_EVNTSEL3       | Package | Uncore C-Box 0 Perfmon Event Select for C-Box 0 Counter 3 |



Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| E05H             |     | MSR_CO_PMON_BOX_FILTER0    | Package | Uncore C-Box 0 Perfmon Box Wide Filter 0                  |
| E06H             |     | MSR_CO_PMON_BOX_FILTER1    | Package | Uncore C-Box 0 Perfmon Box Wide Filter 1                  |
| E07H             |     | MSR_CO_PMON_BOX_STATUS     | Package | Uncore C-Box 0 Perfmon Box Wide Status                    |
| E08H             |     | MSR_CO_PMON_CTR0           | Package | Uncore C-Box 0 Perfmon Counter 0                          |
| E09H             |     | MSR_CO_PMON_CTR1           | Package | Uncore C-Box 0 Perfmon Counter 1                          |
| E0AH             |     | MSR_CO_PMON_CTR2           | Package | Uncore C-Box 0 Perfmon Counter 2                          |
| E0BH             |     | MSR_CO_PMON_CTR3           | Package | Uncore C-Box 0 Perfmon Counter 3                          |
| E10H             |     | MSR_C1_PMON_BOX_CTL        | Package | Uncore C-Box 1 Perfmon for Box-Wide Control               |
| E11H             |     | MSR_C1_PMON_EVNTSELO       | Package | Uncore C-Box 1 Perfmon Event Select for C-Box 1 Counter 0 |
| E12H             |     | MSR_C1_PMON_EVNTSEL1       | Package | Uncore C-Box 1 Perfmon Event Select for C-Box 1 Counter 1 |
| E13H             |     | MSR_C1_PMON_EVNTSEL2       | Package | Uncore C-Box 1 Perfmon Event Select for C-Box 1 Counter 2 |
| E14H             |     | MSR_C1_PMON_EVNTSEL3       | Package | Uncore C-Box 1 Perfmon Event Select for C-Box 1 Counter 3 |
| E15H             |     | MSR_C1_PMON_BOX_FILTER0    | Package | Uncore C-Box 1 Perfmon Box Wide Filter 0                  |
| E16H             |     | MSR_C1_PMON_BOX_FILTER1    | Package | Uncore C-Box 1 Perfmon Box Wide Filter1                   |
| E17H             |     | MSR_C1_PMON_BOX_STATUS     | Package | Uncore C-Box 1 Perfmon Box Wide Status                    |
| E18H             |     | MSR_C1_PMON_CTR0           | Package | Uncore C-Box 1 Perfmon Counter 0                          |
| E19H             |     | MSR_C1_PMON_CTR1           | Package | Uncore C-Box 1 Perfmon Counter 1                          |
| E1AH             |     | MSR_C1_PMON_CTR2           | Package | Uncore C-Box 1 Perfmon Counter 2                          |
| E1BH             |     | MSR_C1_PMON_CTR3           | Package | Uncore C-Box 1 Perfmon Counter 3                          |
| E20H             |     | MSR_C2_PMON_BOX_CTL        | Package | Uncore C-Box 2 Perfmon for Box-Wide Control               |
| E21H             |     | MSR_C2_PMON_EVNTSELO       | Package | Uncore C-Box 2 Perfmon Event Select for C-Box 2 Counter 0 |
| E22H             |     | MSR_C2_PMON_EVNTSEL1       | Package | Uncore C-Box 2 Perfmon Event Select for C-Box 2 Counter 1 |
| E23H             |     | MSR_C2_PMON_EVNTSEL2       | Package | Uncore C-Box 2 Perfmon Event Select for C-Box 2 Counter 2 |
| E24H             |     | MSR_C2_PMON_EVNTSEL3       | Package | Uncore C-Box 2 Perfmon Event select for C-Box 2 Counter 3 |
| E25H             |     | MSR_C2_PMON_BOX_FILTER0    | Package | Uncore C-Box 2 Perfmon Box Wide Filter 0                  |
| E26H             |     | MSR_C2_PMON_BOX_FILTER1    | Package | Uncore C-Box 2 Perfmon Box Wide Filter1                   |
| E27H             |     | MSR_C2_PMON_BOX_STATUS     | Package | Uncore C-Box 2 Perfmon Box Wide Status                    |
| E28H             |     | MSR_C2_PMON_CTR0           | Package | Uncore C-Box 2 Perfmon Counter 0                          |
| E29H             |     | MSR_C2_PMON_CTR1           | Package | Uncore C-Box 2 Perfmon Counter 1                          |
| E2AH             |     | MSR_C2_PMON_CTR2           | Package | Uncore C-Box 2 Perfmon Counter 2                          |
| E2BH             |     | MSR_C2_PMON_CTR3           | Package | Uncore C-Box 2 Perfmon Counter 3                          |

Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| E30H             |     | MSR_C3_PMON_BOX_CTL        | Package | Uncore C-Box 3 Perfmon for Box-Wide Control               |
| E31H             |     | MSR_C3_PMON_EVNTSELO       | Package | Uncore C-Box 3 Perfmon Event Select for C-Box 3 Counter 0 |
| E32H             |     | MSR_C3_PMON_EVNTSEL1       | Package | Uncore C-Box 3 Perfmon Event Select for C-Box 3 Counter 1 |
| E33H             |     | MSR_C3_PMON_EVNTSEL2       | Package | Uncore C-Box 3 Perfmon Event Select for C-Box 3 Counter 2 |
| E34H             |     | MSR_C3_PMON_EVNTSEL3       | Package | Uncore C-Box 3 Perfmon Event Select for C-Box 3 Counter 3 |
| E35H             |     | MSR_C3_PMON_BOX_FILTER0    | Package | Uncore C-Box 3 Perfmon Box Wide Filter 0                  |
| E36H             |     | MSR_C3_PMON_BOX_FILTER1    | Package | Uncore C-Box 3 Perfmon Box Wide Filter1                   |
| E37H             |     | MSR_C3_PMON_BOX_STATUS     | Package | Uncore C-Box 3 Perfmon Box Wide Status                    |
| E38H             |     | MSR_C3_PMON_CTRL0          | Package | Uncore C-Box 3 Perfmon Counter 0                          |
| E39H             |     | MSR_C3_PMON_CTRL1          | Package | Uncore C-Box 3 Perfmon Counter 1                          |
| E3AH             |     | MSR_C3_PMON_CTRL2          | Package | Uncore C-Box 3 Perfmon Counter 2                          |
| E3BH             |     | MSR_C3_PMON_CTRL3          | Package | Uncore C-Box 3 Perfmon Counter 3                          |
| E40H             |     | MSR_C4_PMON_BOX_CTL        | Package | Uncore C-Box 4 Perfmon for Box-Wide Control               |
| E41H             |     | MSR_C4_PMON_EVNTSELO       | Package | Uncore C-Box 4 Perfmon Event Select for C-Box 4 Counter 0 |
| E42H             |     | MSR_C4_PMON_EVNTSEL1       | Package | Uncore C-Box 4 Perfmon Event Select for C-Box 4 Counter 1 |
| E43H             |     | MSR_C4_PMON_EVNTSEL2       | Package | Uncore C-Box 4 Perfmon Event Select for C-Box 4 Counter 2 |
| E44H             |     | MSR_C4_PMON_EVNTSEL3       | Package | Uncore C-Box 4 Perfmon Event Select for C-Box 4 Counter 3 |
| E45H             |     | MSR_C4_PMON_BOX_FILTER0    | Package | Uncore C-Box 4 Perfmon Box Wide Filter 0                  |
| E46H             |     | MSR_C4_PMON_BOX_FILTER1    | Package | Uncore C-Box 4 Perfmon Box Wide Filter1                   |
| E47H             |     | MSR_C4_PMON_BOX_STATUS     | Package | Uncore C-Box 4 Perfmon Box Wide Status                    |
| E48H             |     | MSR_C4_PMON_CTRL0          | Package | Uncore C-Box 4 Perfmon Counter 0                          |
| E49H             |     | MSR_C4_PMON_CTRL1          | Package | Uncore C-Box 4 Perfmon Counter 1                          |
| E4AH             |     | MSR_C4_PMON_CTRL2          | Package | Uncore C-Box 4 Perfmon Counter 2                          |
| E4BH             |     | MSR_C4_PMON_CTRL3          | Package | Uncore C-Box 4 Perfmon Counter 3                          |
| E50H             |     | MSR_C5_PMON_BOX_CTL        | Package | Uncore C-Box 5 Perfmon for Box-Wide Control               |
| E51H             |     | MSR_C5_PMON_EVNTSELO       | Package | Uncore C-Box 5 Perfmon Event Select for C-Box 5 Counter 0 |
| E52H             |     | MSR_C5_PMON_EVNTSEL1       | Package | Uncore C-Box 5 Perfmon Event Select for C-Box 5 Counter 1 |
| E53H             |     | MSR_C5_PMON_EVNTSEL2       | Package | Uncore C-Box 5 Perfmon Event Select for C-Box 5 Counter 2 |
| E54H             |     | MSR_C5_PMON_EVNTSEL3       | Package | Uncore C-Box 5 Perfmon Event Select for C-Box 5 Counter 3 |

Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| E55H             |     | MSR_C5_PMON_BOX_FILTER0    | Package | Uncore C-Box 5 Perfmon Box Wide Filter 0                  |
| E56H             |     | MSR_C5_PMON_BOX_FILTER1    | Package | Uncore C-Box 5 Perfmon Box Wide Filter 1                  |
| E57H             |     | MSR_C5_PMON_BOX_STATUS     | Package | Uncore C-Box 5 Perfmon Box Wide Status                    |
| E58H             |     | MSR_C5_PMON_CTR0           | Package | Uncore C-Box 5 Perfmon Counter 0                          |
| E59H             |     | MSR_C5_PMON_CTR1           | Package | Uncore C-Box 5 Perfmon Counter 1                          |
| E5AH             |     | MSR_C5_PMON_CTR2           | Package | Uncore C-Box 5 Perfmon Counter 2                          |
| E5BH             |     | MSR_C5_PMON_CTR3           | Package | Uncore C-Box 5 Perfmon Counter 3                          |
| E60H             |     | MSR_C6_PMON_BOX_CTL        | Package | Uncore C-Box 6 Perfmon for Box-Wide Control               |
| E61H             |     | MSR_C6_PMON_EVNTSELO       | Package | Uncore C-Box 6 Perfmon Event Select for C-Box 6 Counter 0 |
| E62H             |     | MSR_C6_PMON_EVNTSEL1       | Package | Uncore C-Box 6 Perfmon Event Select for C-Box 6 Counter 1 |
| E63H             |     | MSR_C6_PMON_EVNTSEL2       | Package | Uncore C-Box 6 Perfmon Event Select for C-Box 6 Counter 2 |
| E64H             |     | MSR_C6_PMON_EVNTSEL3       | Package | Uncore C-Box 6 Perfmon Event Select for C-Box 6 Counter 3 |
| E65H             |     | MSR_C6_PMON_BOX_FILTER0    | Package | Uncore C-Box 6 Perfmon Box Wide Filter 0                  |
| E66H             |     | MSR_C6_PMON_BOX_FILTER1    | Package | Uncore C-Box 6 Perfmon Box Wide Filter 1                  |
| E67H             |     | MSR_C6_PMON_BOX_STATUS     | Package | Uncore C-Box 6 Perfmon Box Wide Status                    |
| E68H             |     | MSR_C6_PMON_CTR0           | Package | Uncore C-Box 6 Perfmon Counter 0                          |
| E69H             |     | MSR_C6_PMON_CTR1           | Package | Uncore C-Box 6 Perfmon Counter 1                          |
| E6AH             |     | MSR_C6_PMON_CTR2           | Package | Uncore C-Box 6 Perfmon Counter 2                          |
| E6BH             |     | MSR_C6_PMON_CTR3           | Package | Uncore C-Box 6 Perfmon Counter 3                          |
| E70H             |     | MSR_C7_PMON_BOX_CTL        | Package | Uncore C-Box 7 Perfmon for Box-Wide Control               |
| E71H             |     | MSR_C7_PMON_EVNTSELO       | Package | Uncore C-Box 7 Perfmon Event Select for C-Box 7 Counter 0 |
| E72H             |     | MSR_C7_PMON_EVNTSEL1       | Package | Uncore C-Box 7 Perfmon Event Select for C-Box 7 Counter 1 |
| E73H             |     | MSR_C7_PMON_EVNTSEL2       | Package | Uncore C-Box 7 Perfmon Event Select for C-Box 7 Counter 2 |
| E74H             |     | MSR_C7_PMON_EVNTSEL3       | Package | Uncore C-Box 7 Perfmon Event Select for C-Box 7 Counter 3 |
| E75H             |     | MSR_C7_PMON_BOX_FILTER0    | Package | Uncore C-Box 7 Perfmon Box Wide Filter 0                  |
| E76H             |     | MSR_C7_PMON_BOX_FILTER1    | Package | Uncore C-Box 7 Perfmon Box Wide Filter 1                  |
| E77H             |     | MSR_C7_PMON_BOX_STATUS     | Package | Uncore C-Box 7 Perfmon Box Wide Status                    |
| E78H             |     | MSR_C7_PMON_CTR0           | Package | Uncore C-Box 7 Perfmon Counter 0                          |
| E79H             |     | MSR_C7_PMON_CTR1           | Package | Uncore C-Box 7 Perfmon Counter 1                          |
| E7AH             |     | MSR_C7_PMON_CTR2           | Package | Uncore C-Box 7 Perfmon Counter 2                          |
| E7BH             |     | MSR_C7_PMON_CTR3           | Package | Uncore C-Box 7 Perfmon Counter 3                          |

**Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| E80H             |     | MSR_C8_PMON_BOX_CTL        | Package | Uncore C-Box 8 Perfmon Local Box Wide Control               |
| E81H             |     | MSR_C8_PMON_EVNTSELO       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 0   |
| E82H             |     | MSR_C8_PMON_EVNTSEL1       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 1   |
| E83H             |     | MSR_C8_PMON_EVNTSEL2       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 2   |
| E84H             |     | MSR_C8_PMON_EVNTSEL3       | Package | Uncore C-Box 8 Perfmon Event Select for C-Box 8 Counter 3   |
| E85H             |     | MSR_C8_PMON_BOX_FILTER0    | Package | Uncore C-Box 8 Perfmon Box Wide Filter 0                    |
| E86H             |     | MSR_C8_PMON_BOX_FILTER1    | Package | Uncore C-Box 8 Perfmon Box Wide Filter 1                    |
| E87H             |     | MSR_C8_PMON_BOX_STATUS     | Package | Uncore C-Box 8 Perfmon Box Wide Status                      |
| E88H             |     | MSR_C8_PMON_CTR0           | Package | Uncore C-Box 8 Perfmon Counter 0                            |
| E89H             |     | MSR_C8_PMON_CTR1           | Package | Uncore C-Box 8 Perfmon Counter 1                            |
| E8AH             |     | MSR_C8_PMON_CTR2           | Package | Uncore C-Box 8 Perfmon Counter 2                            |
| E8BH             |     | MSR_C8_PMON_CTR3           | Package | Uncore C-Box 8 Perfmon Counter 3                            |
| E90H             |     | MSR_C9_PMON_BOX_CTL        | Package | Uncore C-Box 9 Perfmon Local Box Wide Control               |
| E91H             |     | MSR_C9_PMON_EVNTSELO       | Package | Uncore C-Box 9 Perfmon Event Select for C-Box 9 Counter 0   |
| E92H             |     | MSR_C9_PMON_EVNTSEL1       | Package | Uncore C-Box 9 Perfmon Event Select for C-Box 9 Counter 1   |
| E93H             |     | MSR_C9_PMON_EVNTSEL2       | Package | Uncore C-Box 9 Perfmon Event Select for C-Box 9 Counter 2   |
| E94H             |     | MSR_C9_PMON_EVNTSEL3       | Package | Uncore C-Box 9 Perfmon Event Select for C-Box 9 Counter 3   |
| E95H             |     | MSR_C9_PMON_BOX_FILTER0    | Package | Uncore C-Box 9 Perfmon Box Wide Filter 0                    |
| E96H             |     | MSR_C9_PMON_BOX_FILTER1    | Package | Uncore C-Box 9 Perfmon Box Wide Filter 1                    |
| E97H             |     | MSR_C9_PMON_BOX_STATUS     | Package | Uncore C-Box 9 Perfmon Box Wide Status                      |
| E98H             |     | MSR_C9_PMON_CTR0           | Package | Uncore C-Box 9 Perfmon Counter 0                            |
| E99H             |     | MSR_C9_PMON_CTR1           | Package | Uncore C-Box 9 Perfmon Counter 1                            |
| E9AH             |     | MSR_C9_PMON_CTR2           | Package | Uncore C-Box 9 Perfmon Counter 2                            |
| E9BH             |     | MSR_C9_PMON_CTR3           | Package | Uncore C-Box 9 Perfmon Counter 3                            |
| EA0H             |     | MSR_C10_PMON_BOX_CTL       | Package | Uncore C-Box 10 Perfmon Local Box Wide Control              |
| EA1H             |     | MSR_C10_PMON_EVNTSELO      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 0 |
| EA2H             |     | MSR_C10_PMON_EVNTSEL1      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 1 |
| EA3H             |     | MSR_C10_PMON_EVNTSEL2      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 2 |
| EA4H             |     | MSR_C10_PMON_EVNTSEL3      | Package | Uncore C-Box 10 Perfmon Event Select for C-Box 10 Counter 3 |

Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| EA5H             |     | MSR_C10_PMON_BOX_FILTER0   | Package | Uncore C-Box 10 Perfmon Box Wide Filter 0                   |
| EA6H             |     | MSR_C10_PMON_BOX_FILTER1   | Package | Uncore C-Box 10 Perfmon Box Wide Filter 1                   |
| EA7H             |     | MSR_C10_PMON_BOX_STATUS    | Package | Uncore C-Box 10 Perfmon Box Wide Status                     |
| EA8H             |     | MSR_C10_PMON_CTR0          | Package | Uncore C-Box 10 Perfmon Counter 0                           |
| EA9H             |     | MSR_C10_PMON_CTR1          | Package | Uncore C-Box 10 Perfmon Counter 1                           |
| EA AH            |     | MSR_C10_PMON_CTR2          | Package | Uncore C-Box 10 Perfmon Counter 2                           |
| EABH             |     | MSR_C10_PMON_CTR3          | Package | Uncore C-Box 10 Perfmon Counter 3                           |
| EBOH             |     | MSR_C11_PMON_BOX_CTL       | Package | Uncore C-Box 11 Perfmon Local Box Wide Control              |
| EB1H             |     | MSR_C11_PMON_EVNTSELO      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 0 |
| EB2H             |     | MSR_C11_PMON_EVNTSEL1      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 1 |
| EB3H             |     | MSR_C11_PMON_EVNTSEL2      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 2 |
| EB4H             |     | MSR_C11_PMON_EVNTSEL3      | Package | Uncore C-Box 11 Perfmon Event Select for C-Box 11 Counter 3 |
| EB5H             |     | MSR_C11_PMON_BOX_FILTER0   | Package | Uncore C-Box 11 Perfmon Box Wide Filter 0                   |
| EB6H             |     | MSR_C11_PMON_BOX_FILTER1   | Package | Uncore C-Box 11 Perfmon Box Wide Filter 1                   |
| EB7H             |     | MSR_C11_PMON_BOX_STATUS    | Package | Uncore C-Box 11 Perfmon Box Wide Status                     |
| EB8H             |     | MSR_C11_PMON_CTR0          | Package | Uncore C-Box 11 Perfmon Counter 0                           |
| EB9H             |     | MSR_C11_PMON_CTR1          | Package | Uncore C-Box 11 Perfmon Counter 1                           |
| EBAH             |     | MSR_C11_PMON_CTR2          | Package | Uncore C-Box 11 Perfmon Counter 2                           |
| EBBH             |     | MSR_C11_PMON_CTR3          | Package | Uncore C-Box 11 Perfmon Counter 3                           |
| EC0H             |     | MSR_C12_PMON_BOX_CTL       | Package | Uncore C-Box 12 Perfmon Local Box Wide Control              |
| EC1H             |     | MSR_C12_PMON_EVNTSELO      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 0 |
| EC2H             |     | MSR_C12_PMON_EVNTSEL1      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 1 |
| EC3H             |     | MSR_C12_PMON_EVNTSEL2      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 2 |
| EC4H             |     | MSR_C12_PMON_EVNTSEL3      | Package | Uncore C-Box 12 Perfmon Event Select for C-Box 12 Counter 3 |
| EC5H             |     | MSR_C12_PMON_BOX_FILTER0   | Package | Uncore C-Box 12 Perfmon Box Wide Filter 0                   |
| EC6H             |     | MSR_C12_PMON_BOX_FILTER1   | Package | Uncore C-Box 12 Perfmon Box Wide Filter 1                   |
| EC7H             |     | MSR_C12_PMON_BOX_STATUS    | Package | Uncore C-Box 12 Perfmon Box Wide Status                     |
| EC8H             |     | MSR_C12_PMON_CTR0          | Package | Uncore C-Box 12 Perfmon Counter 0                           |
| EC9H             |     | MSR_C12_PMON_CTR1          | Package | Uncore C-Box 12 Perfmon Counter 1                           |
| ECAH             |     | MSR_C12_PMON_CTR2          | Package | Uncore C-Box 12 Perfmon Counter 2                           |
| ECBH             |     | MSR_C12_PMON_CTR3          | Package | Uncore C-Box 12 Perfmon Counter 3                           |

**Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| ED0H             |     | MSR_C13_PMON_BOX_CTL       | Package | Uncore C-Box 13 Perfmon local box wide control.             |
| ED1H             |     | MSR_C13_PMON_EVNTSELO      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 0 |
| ED2H             |     | MSR_C13_PMON_EVNTSEL1      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 1 |
| ED3H             |     | MSR_C13_PMON_EVNTSEL2      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 2 |
| ED4H             |     | MSR_C13_PMON_EVNTSEL3      | Package | Uncore C-Box 13 Perfmon Event Select for C-Box 13 Counter 3 |
| ED5H             |     | MSR_C13_PMON_BOX_FILTER0   | Package | Uncore C-Box 13 Perfmon Box Wide Filter 0                   |
| ED6H             |     | MSR_C13_PMON_BOX_FILTER1   | Package | Uncore C-Box 13 Perfmon Box Wide Filter 1                   |
| ED7H             |     | MSR_C13_PMON_BOX_STATUS    | Package | Uncore C-Box 13 Perfmon Box Wide Status                     |
| ED8H             |     | MSR_C13_PMON_CTR0          | Package | Uncore C-Box 13 Perfmon Counter 0                           |
| ED9H             |     | MSR_C13_PMON_CTR1          | Package | Uncore C-Box 13 Perfmon Counter 1                           |
| EDAH             |     | MSR_C13_PMON_CTR2          | Package | Uncore C-Box 13 Perfmon Counter 2                           |
| EDBH             |     | MSR_C13_PMON_CTR3          | Package | Uncore C-Box 13 Perfmon Counter 3                           |
| EE0H             |     | MSR_C14_PMON_BOX_CTL       | Package | Uncore C-Box 14 Perfmon Local Box Wide Control              |
| EE1H             |     | MSR_C14_PMON_EVNTSELO      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 0 |
| EE2H             |     | MSR_C14_PMON_EVNTSEL1      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 1 |
| EE3H             |     | MSR_C14_PMON_EVNTSEL2      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 2 |
| EE4H             |     | MSR_C14_PMON_EVNTSEL3      | Package | Uncore C-Box 14 Perfmon Event Select for C-Box 14 Counter 3 |
| EE5H             |     | MSR_C14_PMON_BOX_FILTER    | Package | Uncore C-Box 14 Perfmon Box Wide Filter 0                   |
| EE6H             |     | MSR_C14_PMON_BOX_FILTER1   | Package | Uncore C-Box 14 Perfmon Box Wide Filter 1                   |
| EE7H             |     | MSR_C14_PMON_BOX_STATUS    | Package | Uncore C-Box 14 Perfmon Box Wide Status                     |
| EE8H             |     | MSR_C14_PMON_CTR0          | Package | Uncore C-Box 14 Perfmon Counter 0                           |
| EE9H             |     | MSR_C14_PMON_CTR1          | Package | Uncore C-Box 14 Perfmon Counter 1                           |
| EEAH             |     | MSR_C14_PMON_CTR2          | Package | Uncore C-Box 14 Perfmon Counter 2                           |
| EEBH             |     | MSR_C14_PMON_CTR3          | Package | Uncore C-Box 14 Perfmon Counter 3                           |
| EF0H             |     | MSR_C15_PMON_BOX_CTL       | Package | Uncore C-Box 15 Perfmon Local Box Wide Control              |
| EF1H             |     | MSR_C15_PMON_EVNTSELO      | Package | Uncore C-Box 15 Perfmon Event Select for C-Box 15 Counter 0 |
| EF2H             |     | MSR_C15_PMON_EVNTSEL1      | Package | Uncore C-Box 15 Perfmon Event Select for C-Box 15 Counter 1 |
| EF3H             |     | MSR_C15_PMON_EVNTSEL2      | Package | Uncore C-Box 15 Perfmon Event Select for C-Box 15 Counter 2 |
| EF4H             |     | MSR_C15_PMON_EVNTSEL3      | Package | Uncore C-Box 15 Perfmon Event Select for C-Box 15 Counter 3 |

Table 2-32. Uncore PMU MSRs in Intel® Xeon® Processor E5 v3 Family (Contd.)

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| EF5H             |     | MSR_C15_PMON_BOX_FILTER0   | Package | Uncore C-Box 15 Perfmon Box Wide Filter 0                   |
| EF6H             |     | MSR_C15_PMON_BOX_FILTER1   | Package | Uncore C-Box 15 Perfmon Box Wide Filter 1                   |
| EF7H             |     | MSR_C15_PMON_BOX_STATUS    | Package | Uncore C-Box 15 Perfmon Box Wide Status                     |
| EF8H             |     | MSR_C15_PMON_CTR0          | Package | Uncore C-Box 15 Perfmon Counter 0                           |
| EF9H             |     | MSR_C15_PMON_CTR1          | Package | Uncore C-Box 15 Perfmon Counter 1                           |
| EFAH             |     | MSR_C15_PMON_CTR2          | Package | Uncore C-Box 15 Perfmon Counter 2                           |
| EFBH             |     | MSR_C15_PMON_CTR3          | Package | Uncore C-Box 15 Perfmon Counter 3                           |
| F00H             |     | MSR_C16_PMON_BOX_CTL       | Package | Uncore C-Box 16 Perfmon for Box-Wide Control                |
| F01H             |     | MSR_C16_PMON_EVNTSELO      | Package | Uncore C-Box 16 Perfmon Event Select for C-Box 16 Counter 0 |
| F02H             |     | MSR_C16_PMON_EVNTSEL1      | Package | Uncore C-Box 16 Perfmon Event Select for C-Box 16 Counter 1 |
| F03H             |     | MSR_C16_PMON_EVNTSEL2      | Package | Uncore C-Box 16 Perfmon Event Select for C-Box 16 Counter 2 |
| F04H             |     | MSR_C16_PMON_EVNTSEL3      | Package | Uncore C-Box 16 Perfmon Event Select for C-Box 16 Counter 3 |
| F05H             |     | MSR_C16_PMON_BOX_FILTER0   | Package | Uncore C-Box 16 Perfmon Box Wide Filter 0                   |
| F06H             |     | MSR_C16_PMON_BOX_FILTER1   | Package | Uncore C-Box 16 Perfmon Box Wide Filter 1                   |
| F07H             |     | MSR_C16_PMON_BOX_STATUS    | Package | Uncore C-Box 16 Perfmon Box Wide Status                     |
| F08H             |     | MSR_C16_PMON_CTR0          | Package | Uncore C-Box 16 Perfmon Counter 0                           |
| F09H             |     | MSR_C16_PMON_CTR1          | Package | Uncore C-Box 16 Perfmon Counter 1                           |
| F0AH             |     | MSR_C16_PMON_CTR2          | Package | Uncore C-Box 16 Perfmon Counter 2                           |
| E0BH             |     | MSR_C16_PMON_CTR3          | Package | Uncore C-Box 16 Perfmon Counter 3                           |
| F10H             |     | MSR_C17_PMON_BOX_CTL       | Package | Uncore C-Box 17 Perfmon for Box-Wide Control                |
| F11H             |     | MSR_C17_PMON_EVNTSELO      | Package | Uncore C-Box 17 Perfmon Event Select for C-Box 17 Counter 0 |
| F12H             |     | MSR_C17_PMON_EVNTSEL1      | Package | Uncore C-Box 17 Perfmon Event Select for C-Box 17 Counter 1 |
| F13H             |     | MSR_C17_PMON_EVNTSEL2      | Package | Uncore C-Box 17 Perfmon Event Select for C-Box 17 Counter 2 |
| F14H             |     | MSR_C17_PMON_EVNTSEL3      | Package | Uncore C-Box 17 Perfmon Event Select for C-Box 17 Counter 3 |
| F15H             |     | MSR_C17_PMON_BOX_FILTER0   | Package | Uncore C-Box 17 Perfmon Box Wide Filter 0                   |
| F16H             |     | MSR_C17_PMON_BOX_FILTER1   | Package | Uncore C-Box 17 Perfmon Box Wide Filter1                    |
| F17H             |     | MSR_C17_PMON_BOX_STATUS    | Package | Uncore C-Box 17 Perfmon Box Wide Status                     |
| F18H             |     | MSR_C17_PMON_CTR0          | Package | Uncore C-Box 17 Perfmon Counter 0                           |
| F19H             |     | MSR_C17_PMON_CTR1          | Package | Uncore C-Box 17 Perfmon Counter 1                           |
| F1AH             |     | MSR_C17_PMON_CTR2          | Package | Uncore C-Box 17 Perfmon Counter 2                           |
| F1BH             |     | MSR_C17_PMON_CTR3          | Package | Uncore C-Box 17 Perfmon Counter 3                           |

## 2.14 MSRS IN INTEL® CORE™ M PROCESSORS AND 5TH GENERATION INTEL CORE PROCESSORS

The Intel® Core™ M-5xxx processors and 5th generation Intel® Core™ Processors, and Intel® Xeon® Processor E3-1200 v4 family are based on the Broadwell microarchitecture. The Intel® Core™ M-5xxx processors and 5th generation Intel® Core™ Processors have CPUID DisplayFamily\_DisplayModel signature 06\_3DH. Intel® Xeon® Processor E3-1200 v4 family and the 5th generation Intel® Core™ Processors have CPUID DisplayFamily\_DisplayModel signature 06\_47H. Processors with signatures 06\_3DH and 06\_47H support the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-21, Table 2-24, Table 2-28, Table 2-29, Table 2-33, and Table 2-34. For an MSR listed in Table 2-34 that also appears in the model-specific tables of prior generations, Table 2-34 supercede prior generation tables.

Table 2-33 lists MSRs that are common to processors based on the Broadwell microarchitectures (including CPUID signatures 06\_3DH, 06\_47H, 06\_4FH, and 06\_56H).

**Table 2-33. Additional MSRs Common to Processors Based the Broadwell Microarchitectures**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
| 38EH             | 910 | IA32_PERF_GLOBAL_STATUS    | Thread | See Table 2-2. See Section 18.6.2.2, “Global Counter Control Facilities.”     |
|                  |     | 0                          |        | Ovf_PMC0  |
|                  |     | 1                          |        | Ovf_PMC1  |
|                  |     | 2                          |        | Ovf_PMC2  |
|                  |     | 3                          |        | Ovf_PMC3  |
|                  |     | 31:4                       |        | Reserved  |
|                  |     | 32                         |        | Ovf_FixedCtr0   |
|                  |     | 33                         |        | Ovf_FixedCtr1   |
|                  |     | 34                         |        | Ovf_FixedCtr2   |
|                  |     | 54:35                      |        | Reserved  |
|                  |     | 55                         |        | Trace_ToPA_PMI<br>See Section 35.2.6.2, “Table of Physical Addresses (ToPA).” |
|                  |     | 60:56                      |        | Reserved  |
|                  |     | 61                         |        | Ovf_Uncore  |
|                  |     | 62                         |        | Ovf_BufDSSAVE   |
| 63               |     | CondChgd                   |        |   |
| 390H             | 912 | IA32_PERF_GLOBAL_OVF_CTRL  | Thread | See Table 2-2. See Section 18.6.2.2, “Global Counter Control Facilities.”     |
|                  |     | 0                          |        | Set 1 to clear Ovf_PMC0   |
|                  |     | 1                          |        | Set 1 to clear Ovf_PMC1   |
|                  |     | 2                          |        | Set 1 to clear Ovf_PMC2   |
|                  |     | 3                          |        | Set 1 to clear Ovf_PMC3   |
|                  |     | 31:4                       |        | Reserved  |
|                  |     | 32                         |        | Set 1 to clear Ovf_FixedCtr0  |
|                  |     | 33                         |        | Set 1 to clear Ovf_FixedCtr1  |



Table 2-33. Additional MSRs Common to Processors Based the Broadwell Microarchitectures

| Register Address |      | Register Name / Bit Fields   | Scope  | Bit Description   |
|------------------|------|------------------------------|--------|---|
| Hex              | Dec  |                              |        |   |
|                  |      | 34                           |        | Set 1 to clear Ovf_FixedCtr2  |
|                  |      | 54:35                        |        | Reserved.   |
|                  |      | 55                           |        | Set 1 to clear Trace_ToPA_PMI.<br>See Section 35.2.6.2, "Table of Physical Addresses (ToPA)." |
|                  |      | 60:56                        |        | Reserved  |
|                  |      | 61                           |        | Set 1 to clear Ovf_Uncore   |
|                  |      | 62                           |        | Set 1 to clear Ovf_BufDSSAVE  |
|                  |      | 63                           |        | Set 1 to clear CondChgd   |
| 560H             | 1376 | IA32_RTIT_OUTPUT_BASE        | THREAD | Trace Output Base Register (R/W)  |
|                  |      | 6:0                          |        | Reserved  |
|                  |      | MAXPHYADDR <sup>1</sup> -1:7 |        | Base physical address.  |
|                  |      | 63:MAXPHYADDR                |        | Reserved  |
| 561H             | 1377 | IA32_RTIT_OUTPUT_MASK_PTRS   | THREAD | Trace Output Mask Pointers Register (R/W)   |
|                  |      | 6:0                          |        | Reserved  |
|                  |      | 31:7                         |        | MaskOrTableOffset   |
|                  |      | 63:32                        |        | Output Offset.  |
| 570H             | 1392 | IA32_RTIT_CTL                | Thread | Trace Control Register (R/W)  |
|                  |      | 0                            |        | TraceEn   |
|                  |      | 1                            |        | Reserved, must be zero.   |
|                  |      | 2                            |        | OS  |
|                  |      | 3                            |        | User  |
|                  |      | 6:4                          |        | Reserved, must be zero.   |
|                  |      | 7                            |        | CR3 filter  |
|                  |      | 8                            |        | ToPA<br>Writing 0 will #GP if also setting TraceEn.   |
|                  |      | 9                            |        | Reserved, must be zero.   |
|                  |      | 10                           |        | TSCEn   |
|                  |      | 11                           |        | DisRETC   |
|                  |      | 12                           |        | Reserved, must be zero.   |
|                  |      | 13                           |        | Reserved; writing 0 will #GP if also setting TraceEn.   |
|                  |      | 63:14                        |        | Reserved, must be zero.   |
| 571H             | 1393 | IA32_RTIT_STATUS             | Thread | Tracing Status Register (R/W)   |
|                  |      | 0                            |        | Reserved, writes ignored.   |
|                  |      | 1                            |        | ContexEn, writes ignored.   |
|                  |      | 2                            |        | TriggerEn, writes ignored.  |
|                  |      | 3                            |        | Reserved  |

**Table 2-33. Additional MSRs Common to Processors Based the Broadwell Microarchitectures**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 4                          |         | Error (R/W)   |
|                  |      | 5                          |         | Stopped   |
|                  |      | 63:6                       |         | Reserved, must be zero.   |
| 572H             | 1394 | IA32_RTIT_CR3_MATCH        | THREAD  | Trace Filter CR3 Match Register (R/W)   |
|                  |      | 4:0                        |         | Reserved  |
|                  |      | 63:5                       |         | CR3[63:5] value to match.   |
| 620H             |      | MSR UNCORE_RATIO_LIMIT     | Package | Uncore Ratio Limit (R/W)<br>Out of reset, the min_ratio and max_ratio fields represent the widest possible range of uncore frequencies. Writing to these fields allows software to control the minimum and the maximum frequency that hardware will select. |
|                  |      | 63:15                      |         | Reserved  |
|                  |      | 14:8                       |         | MIN_RATIO<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.   |
|                  |      | 7                          |         | Reserved  |
|                  |      | 6:0                        |         | MAX_RATIO<br>This field is used to limit the max ratio of the LLC/Ring.   |

**NOTES:**

1. MAXPHYADDR is reported by CPUID.80000008H:EAX[7:0].

Table 2-34 lists MSRs that are specific to Intel Core M processors and 5th Generation Intel Core Processors.

**Table 2-34. Additional MSRs Supported by Intel® Core™ M Processors and 5th Generation Intel® Core™ Processors**

| Register Address |     | Register Name              | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core  | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> . |

Table 2-34. Additional MSRs Supported by Intel® Core™ M Processors and 5th Generation Intel® Core™ Processors

| Register Address |     | Register Name         | Scope   | Bit Description  |
|------------------|-----|-----------------------|---------|--|
| Hex              | Dec |                       |         |  |
|                  |     | 3:0                   |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>0000b: C0/C1 (no package C-state support)<br>0001b: C2<br>0010b: C3<br>0011b: C6<br>0100b: C7<br>0101b: C7s<br>0110b: C8<br>0111b: C9<br>1000b: C10 |
|                  |     | 9:4                   |         | Reserved   |
|                  |     | 10                    |         | I/O MWAIT Redirection Enable (R/W)   |
|                  |     | 14:11                 |         | Reserved   |
|                  |     | 15                    |         | CFG Lock (R/W0)  |
|                  |     | 24:16                 |         | Reserved   |
|                  |     | 25                    |         | C3 State Auto Demotion Enable (R/W)  |
|                  |     | 26                    |         | C1 State Auto Demotion Enable (R/W)  |
|                  |     | 27                    |         | Enable C3 Undemotion (R/W)   |
|                  |     | 28                    |         | Enable C1 Undemotion (R/W)   |
|                  |     | 29                    |         | Enable Package C-State Auto-Demotion (R/W)   |
|                  |     | 30                    |         | Enable Package C-State Undemotion (R/W)  |
|                  |     | 63:31                 |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |     | 7:0                   | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.  |
|                  |     | 15:8                  | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.  |
|                  |     | 23:16                 | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.  |
|                  |     | 31:24                 | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.  |

**Table 2-34. Additional MSRs Supported by Intel® Core™ M Processors and 5th Generation Intel® Core™ Processors**

| Register Address |      | Register Name         | Scope   | Bit Description  |
|------------------|------|-----------------------|---------|--|
| Hex              | Dec  |                       |         |  |
|                  |      | 39:32                 | Package | Maximum Ratio Limit for 5C<br>Maximum turbo ratio limit of 5core active. |
|                  |      | 47:40                 | Package | Maximum Ratio Limit for 6C<br>Maximum turbo ratio limit of 6core active. |
|                  |      | 63:48                 |         | Reserved   |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."   |

See Table 2-19, Table 2-20, Table 2-21, Table 2-24, Table 2-28, Table 2-29, Table 2-33 for other MSR definitions applicable to processors with CPUID signature 06\_3DH.

## 2.15 MSRS IN INTEL® XEON® PROCESSORS E5 V4 FAMILY

The MSRs listed in Table 2-35 are available and common to Intel® Xeon® Processor D product Family (CPUID DisplayFamily\_DisplayModel = 06\_56H) and to Intel Xeon processors E5 v4, E7 v4 families (CPUID DisplayFamily\_DisplayModel = 06\_4FH). They are based on the Broadwell microarchitecture.

See Section 2.15.1 for lists of tables of MSRs that are supported by Intel® Xeon® Processor D Family.

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 4EH              | 78  | MSR_PPIN_CTL               | Package | Protected Processor Inventory Number Enable Control (R/W)  |
|                  |     | 0                          |         | LockOut (R/W/O)<br>See Table 2-25.   |
|                  |     | 1                          |         | Enable_PPIN (R/W)<br>See Table 2-25.   |
|                  |     | 63:2                       |         | Reserved   |
| 4FH              | 79  | MSR_PPIN                   | Package | Protected Processor Inventory Number (R/O)   |
|                  |     | 63:0                       |         | Protected Processor Inventory Number (R/O)<br>See Table 2-25.  |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> . |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>See Table 2-25.   |
|                  |     | 22:16                      |         | Reserved.  |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 23                         | Package | PPIN_CAP (R/O)<br>See Table 2-25.  |
|                  |     | 27:24                      |         | Reserved   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>See Table 2-25.   |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>See Table 2-25.   |
|                  |     | 30                         | Package | Programmable TJ OFFSET (R/O)<br>See Table 2-25.  |
|                  |     | 39:31                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>See Table 2-25.  |
|                  |     | 63:48                      |         | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-state support)<br>001b: C2<br>010b: C6 (non-retention)<br>011b: C6 (retention)<br>111b: No Package C state limits. All C states supported by the processor are available. |
|                  |     | 9:3                        |         | Reserved   |
|                  |     | 10                         |         | I/O MWAIT Redirection Enable (R/W)   |
|                  |     | 14:11                      |         | Reserved   |
|                  |     | 15                         |         | CFG Lock (R/WO)  |
|                  |     | 16                         |         | Automatic C-State Conversion Enable (R/W)<br>If 1, the processor will convert HALT or MWAIT(C1) to MWAIT(C6).  |
|                  |     | 24:17                      |         | Reserved   |
|                  |     | 25                         |         | C3 State Auto Demotion Enable (R/W)  |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
|                  |     | 26                         |        | C1 State Auto Demotion Enable (R/W)   |
|                  |     | 27                         |        | Enable C3 Undemotion (R/W)  |
|                  |     | 28                         |        | Enable C1 Undemotion (R/W)  |
|                  |     | 29                         |        | Package C State Demotion Enable (R/W)   |
|                  |     | 30                         |        | Package C State UnDemotion Enable (R/W)   |
|                  |     | 63:31                      |        | Reserved  |
| 179H             | 377 | IA32_MCG_CAP               | Thread | Global Machine Check Capability (R/O)   |
|                  |     | 7:0                        |        | Count   |
|                  |     | 8                          |        | MCG_CTL_P   |
|                  |     | 9                          |        | MCG_EXT_P   |
|                  |     | 10                         |        | MCP_CMCI_P  |
|                  |     | 11                         |        | MCG_TES_P   |
|                  |     | 15:12                      |        | Reserved  |
|                  |     | 23:16                      |        | MCG_EXT_CNT   |
|                  |     | 24                         |        | MCG_SER_P   |
|                  |     | 25                         |        | MCG_EM_P  |
|                  |     | 26                         |        | MCG_ELOG_P  |
|                  |     | 63:27                      |        | Reserved  |
| 17DH             | 390 | MSR_SMM_MCA_CAP            | Thread | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability Enhancement. Accessible only while in SMM.   |
|                  |     | 57:0                       |        | Reserved  |
|                  |     | 58                         |        | SMM_Code_Access_Chk (SMM-RO)<br>If set to 1, indicates that the SMM code access restriction is supported and a host-space interface available to SMM handler. |
|                  |     | 59                         |        | Long_Flow_Indication (SMM-RO)<br>If set to 1, indicates that the SMM long flow indicator is supported and a host-space interface available to SMM handler.    |
|                  |     | 63:60                      |        | Reserved  |
| 19CH             | 412 | IA32_THERM_STATUS          | Core   | Thermal Monitor Status (R/W)<br>See Table 2-2.  |
|                  |     | 0                          |        | Thermal Status (RO)<br>See Table 2-2.   |
|                  |     | 1                          |        | Thermal Status Log (R/WCO)<br>See Table 2-2.  |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields           | Scope   | Bit Description   |
|------------------|-----|--------------------------------------|---------|---|
| Hex              | Dec |                                      |         |   |
|                  |     | 2                                    |         | PROTCHOT # or FORCEPR# Status (RO)<br>See Table 2-2.      |
|                  |     | 3                                    |         | PROTCHOT # or FORCEPR# Log (R/WCO)<br>See Table 2-2.      |
|                  |     | 4                                    |         | Critical Temperature Status (RO)<br>See Table 2-2.        |
|                  |     | 5                                    |         | Critical Temperature Status Log (R/WCO)<br>See Table 2-2. |
|                  |     | 6                                    |         | Thermal Threshold #1 Status (RO)<br>See Table 2-2.        |
|                  |     | 7                                    |         | Thermal Threshold #1 Log (R/WCO)<br>See Table 2-2.        |
|                  |     | 8                                    |         | Thermal Threshold #2 Status (RO)<br>See Table 2-2.        |
|                  |     | 9                                    |         | Thermal Threshold #2 Log (R/WCO)<br>See Table 2-2.        |
|                  |     | 10                                   |         | Power Limitation Status (RO)<br>See Table 2-2.            |
|                  |     | 11                                   |         | Power Limitation Log (R/WCO)<br>See Table 2-2.            |
|                  |     | 12                                   |         | Current Limit Status (RO)<br>See Table 2-2.               |
|                  |     | 13                                   |         | Current Limit Log (R/WCO)<br>See Table 2-2.               |
|                  |     | 14                                   |         | Cross Domain Limit Status (RO)<br>See Table 2-2.          |
|                  |     | 15                                   |         | Cross Domain Limit Log (R/WCO)<br>See Table 2-2.          |
|                  |     | 22:16                                |         | Digital Readout (RO)<br>See Table 2-2.                    |
|                  |     | 26:23                                |         | Reserved  |
|                  |     | 30:27                                |         | Resolution in Degrees Celsius (RO)<br>See Table 2-2.      |
| 31               |     | Reading Valid (RO)<br>See Table 2-2. |         |   |
| 63:32            |     | Reserved                             |         |   |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET               | Package | Temperature Target  |
|                  |     | 15:0                                 |         | Reserved  |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 23:16                      |         | Temperature Target (RO)<br>See Table 2-25.   |
|                  |      | 27:24                      |         | TCC Activation Offset (R/W)<br>See Table 2-25.   |
|                  |      | 63:28                      |         | Reserved.  |
| 1ADH             | 429  | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |      | 7:0                        | Package | Maximum Ratio Limit for 1C   |
|                  |      | 15:8                       | Package | Maximum Ratio Limit for 2C   |
|                  |      | 23:16                      | Package | Maximum Ratio Limit for 3C   |
|                  |      | 31:24                      | Package | Maximum Ratio Limit for 4C   |
|                  |      | 39:32                      | Package | Maximum Ratio Limit for 5C   |
|                  |      | 47:40                      | Package | Maximum Ratio Limit for 6C   |
|                  |      | 55:48                      | Package | Maximum Ratio Limit for 7C   |
|                  |      | 63:56                      | Package | Maximum Ratio Limit for 8C   |
| 1AEH             | 430  | MSR_TURBO_RATIO_LIMIT1     | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |      | 7:0                        | Package | Maximum Ratio Limit for 9C   |
|                  |      | 15:8                       | Package | Maximum Ratio Limit for 10C  |
|                  |      | 23:16                      | Package | Maximum Ratio Limit for 11C  |
|                  |      | 31:24                      | Package | Maximum Ratio Limit for 12C  |
|                  |      | 39:32                      | Package | Maximum Ratio Limit for 13C  |
|                  |      | 47:40                      | Package | Maximum Ratio Limit for 14C  |
|                  |      | 55:48                      | Package | Maximum Ratio Limit for 15C  |
|                  |      | 63:56                      | Package | Maximum Ratio Limit for 16C  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers Used in RAPL Interfaces (R/O)   |
|                  |      | 3:0                        | Package | Power Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 7:4                        | Package | Reserved   |
|                  |      | 12:8                       | Package | Energy Status Units<br>Energy related information (in Joules) is based on the multiplier, $1/2^{\text{ESU}}$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 0EH (or 61 micro-joules). |
|                  |      | 15:13                      | Package | Reserved   |



**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description   |
|------------------|------|-----------------------------|---------|---|
| Hex              | Dec  |                             |         |   |
|                  |      | 19:16                       | Package | Time Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 63:20                       |         | Reserved  |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT        | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS      | Package | DRAM Energy Status (R/O)<br>Energy consumed by DRAM devices.  |
|                  |      | 31:0                        |         | Energy in 15.3 micro-joules. Requires BIOS configuration to enable DRAM RAPL mode 0 (Direct VR).  |
|                  |      | 63:32                       |         | Reserved  |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS        | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO         | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 620H             | 1568 | MSR_UNCORE_RATIO_LIMIT      | Package | Uncore Ratio Limit (R/W)<br>Out of reset, the min_ratio and max_ratio fields represent the widest possible range of uncore frequencies. Writing to these fields allows software to control the minimum and the maximum frequency that hardware will select. |
|                  |      | 63:15                       |         | Reserved  |
|                  |      | 14:8                        |         | MIN_RATIO<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.   |
|                  |      | 7                           |         | Reserved  |
|                  |      | 6:0                         |         | MAX_RATIO<br>This field is used to limit the max ratio of the LLC/Ring.   |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS       | Package | Reserved (R/O)<br>Reads return 0.   |
| 690H             | 1680 | MSR_CORE_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.)   |
|                  |      | 0                           |         | PROCHOT Status (R/O)<br>When set, processor core frequency is reduced below the operating system request due to assertion of external PROCHOT.  |
|                  |      | 1                           |         | Thermal Status (R/O)<br>When set, frequency is reduced below the operating system request due to a thermal event.   |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 2                          |       | Power Budget Management Status (R0)<br>When set, frequency is reduced below the operating system request due to PBM limit.  |
|                  |     | 3                          |       | Platform Configuration Services Status (R0)<br>When set, frequency is reduced below the operating system request due to PCS limit.  |
|                  |     | 4                          |       | Reserved  |
|                  |     | 5                          |       | Autonomous Utilization-Based Frequency Control Status (R0)<br>When set, frequency is reduced below the operating system request because the processor has detected that utilization is low.         |
|                  |     | 6                          |       | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.  |
|                  |     | 7                          |       | Reserved  |
|                  |     | 8                          |       | Electrical Design Point Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g., maximum electrical current consumption). |
|                  |     | 9                          |       | Reserved  |
|                  |     | 10                         |       | Multi-Core Turbo Status (R0)<br>When set, frequency is reduced below the operating system request due to Multi-Core Turbo limits.   |
|                  |     | 12:11                      |       | Reserved  |
|                  |     | 13                         |       | Core Frequency P1 Status (R0)<br>When set, frequency is reduced below max non-turbo P1.   |
|                  |     | 14                         |       | Core Max N-Core Turbo Frequency Limiting Status (R0)<br>When set, frequency is reduced below max n-core turbo frequency.  |
|                  |     | 15                         |       | Core Frequency Limiting Status (R0)<br>When set, frequency is reduced below the operating system request.   |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                      |
|                  |     | 18                         |       | Power Budget Management Log<br>When set, indicates that the PBM Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                          |
|                  |     | 19                         |       | Platform Configuration Services Log<br>When set, indicates that the PCS Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                  |
|                  |     | 20                         |       | Reserved  |
|                  |     | 21                         |       | Autonomous Utilization-Based Frequency Control Log<br>When set, indicates that the AUBFC Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |     | 22                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                        |
|                  |     | 23                         |       | Reserved  |
|                  |     | 24                         |       | Electrical Design Point Log<br>When set, indicates that the EDP Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                          |
|                  |     | 25                         |       | Reserved  |
|                  |     | 26                         |       | Multi-Core Turbo Log<br>When set, indicates that the Multi-Core Turbo Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                    |
|                  |     | 28:27                      |       | Reserved  |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 29                         |         | Core Frequency P1 Log<br>When set, indicates that the Core Frequency P1 Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 30                         |         | Core Max N-Core Turbo Frequency Limiting Log<br>When set, indicates that the Core Max n-core Turbo Frequency Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 31                         |         | Core Frequency Limiting Log<br>When set, indicates that the Core Frequency Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                   |
|                  |      | 63:32                      |         | Reserved   |
| 770H             | 1904 | IA32_PM_ENABLE             | Package | See Section 14.4.2, "Enabling HWP".  |
| 771H             | 1905 | IA32_HWP_CAPABILITIES      | Thread  | See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".  |
| 774H             | 1908 | IA32_HWP_REQUEST           | Thread  | See Section 14.4.4, "Managing HWP".  |
|                  |      | 7:0                        |         | Minimum Performance (R/W)  |
|                  |      | 15:8                       |         | Maximum Performance (R/W)  |
|                  |      | 23:16                      |         | Desired Performance (R/W)  |
|                  |      | 63:24                      |         | Reserved   |
| 777H             | 1911 | IA32_HWP_STATUS            | Thread  | See Section 14.4.5, "HWP Feedback".  |
|                  |      | 1:0                        |         | Reserved   |
|                  |      | 2                          |         | Excursion to Minimum (RO)  |
|                  |      | 63:3                       |         | Reserved   |
| C8DH             | 3213 | IA32_QM_EVTSEL             | THREAD  | Monitoring Event Select Register (R/W)<br>If CPUID.(EAX=07H, ECX=0):EBX.RDT-M[bit 12] = 1.   |
|                  |      | 7:0                        |         | EventID (RW)<br>Event encoding:<br>0x00: No monitoring.<br>0x01: L3 occupancy monitoring.<br>0x02: Total memory bandwidth monitoring.<br>0x03: Local memory bandwidth monitoring.<br>All other encoding reserved.                                      |
|                  |      | 31:8                       |         | Reserved   |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 41:32                      |         | RMID (Rw)  |
|                  |      | 63:42                      |         | Reserved   |
| C8FH             | 3215 | IA32_PQR_ASSOC             | THREAD  | Resource Association Register (R/W)  |
|                  |      | 9:0                        |         | RMID   |
|                  |      | 31:10                      |         | Reserved   |
|                  |      | 51:32                      |         | COS (R/W)  |
|                  |      | 63: 52                     |         | Reserved   |
| C90H             | 3216 | IA32_L3_QOS_MASK_0         | Package | L3 Class Of Service Mask - COS 0 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=0. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 0 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C91H             | 3217 | IA32_L3_QOS_MASK_1         | Package | L3 Class Of Service Mask - COS 1 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=1. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 1 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C92H             | 3218 | IA32_L3_QOS_MASK_2         | Package | L3 Class Of Service Mask - COS 2 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=2. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 2 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C93H             | 3219 | IA32_L3_QOS_MASK_3         | Package | L3 Class Of Service Mask - COS 3 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=3. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 3 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C94H             | 3220 | IA32_L3_QOS_MASK_4         | Package | L3 Class Of Service Mask - COS 4 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=4. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 4 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C95H             | 3221 | IA32_L3_QOS_MASK_5         | Package | L3 Class Of Service Mask - COS 5 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=5. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 5 enforcement.                                |
|                  |      | 63:20                      |         | Reserved   |
| C96H             | 3222 | IA32_L3_QOS_MASK_6         | Package | L3 Class Of Service Mask - COS 6 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=6. |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 6 enforcement.                                  |
|                  |      | 63:20                      |         | Reserved   |
| C97H             | 3223 | IA32_L3_QOS_MASK_7         | Package | L3 Class Of Service Mask - COS 7 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=7.   |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 7 enforcement.                                  |
|                  |      | 63:20                      |         | Reserved   |
| C98H             | 3224 | IA32_L3_QOS_MASK_8         | Package | L3 Class Of Service Mask - COS 8 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=8.   |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 8 enforcement.                                  |
|                  |      | 63:20                      |         | Reserved   |
| C99H             | 3225 | IA32_L3_QOS_MASK_9         | Package | L3 Class Of Service Mask - COS 9 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=9.   |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 9 enforcement.                                  |
|                  |      | 63:20                      |         | Reserved   |
| C9AH             | 3226 | IA32_L3_QOS_MASK_10        | Package | L3 Class Of Service Mask - COS 10 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=10. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 10 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9BH             | 3227 | IA32_L3_QOS_MASK_11        | Package | L3 Class Of Service Mask - COS 11 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=11. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 11 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9CH             | 3228 | IA32_L3_QOS_MASK_12        | Package | L3 Class Of Service Mask - COS 12 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=12. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 12 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9DH             | 3229 | IA32_L3_QOS_MASK_13        | Package | L3 Class Of Service Mask - COS 13 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=13. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 13 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9EH             | 3230 | IA32_L3_QOS_MASK_14        | Package | L3 Class Of Service Mask - COS 14 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=14. |

**Table 2-35. Additional MSRs Common to Intel® Xeon® Processor D and Intel Xeon Processors E5 v4 Family Based on the Broadwell Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 14 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9FH             | 3231 | IA32_L3_QOS_MASK_15        | Package | L3 Class Of Service Mask - COS 15 (R/W)<br>If CPUID.(EAX=10H, ECX=1);EDX.COS_MAX[15:0] >=15. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 15 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |

### 2.15.1 Additional MSRs Supported in the Intel® Xeon® Processor D Product Family

The MSRs listed in Table 2-36 are available to Intel® Xeon® Processor D Product Family (CPUID DisplayFamily\_DisplayModel = 06\_56H). The Intel® Xeon® processor D product family is based on the Broadwell microarchitecture and supports the MSR interfaces listed in Table 2-19, Table 2-28, Table 2-33, Table 2-35, and Table 2-36.

**Table 2-36. Additional MSRs Supported by Intel® Xeon® Processor D with DisplayFamily\_DisplayModel 06\_56H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 1ACH             | 428 | MSR_TURBO_RATIO_LIMIT3     | Package | Config Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.   |
|                  |     | 62:0                       | Package | Reserved   |
|                  |     | 63                         | Package | Semaphore for Turbo Ratio Limit Configuration<br>If 1, the processor uses override configuration <sup>1</sup> specified in MSR_TURBO_RATIO_LIMIT, MSR_TURBO_RATIO_LIMIT1.<br>If 0, the processor uses factory-set configuration (Default). |
| 286H             | 646 | IA32_MC6_CTL2              | Package | See Table 2-2.   |
| 287H             | 647 | IA32_MC7_CTL2              | Package | See Table 2-2.   |
| 289H             | 649 | IA32_MC9_CTL2              | Package | See Table 2-2.   |
| 28AH             | 650 | IA32_MC10_CTL2             | Package | See Table 2-2.   |
| 291H             | 657 | IA32_MC17_CTL2             | Package | See Table 2-2.   |
| 292H             | 658 | IA32_MC18_CTL2             | Package | See Table 2-2.   |
| 293H             | 659 | IA32_MC19_CTL2             | Package | See Table 2-2.   |

**Table 2-36. Additional MSRs Supported by Intel® Xeon® Processor D with DisplayFamily\_DisplayModel 06\_56H**

| Register Address   |      | Register Name / Bit Fields | Scope   | Bit Description  |
|--|------|----------------------------|---------|--|
| Hex  | Dec  |                            |         |  |
| 418H   | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC6 reports MC errors from the integrated I/O module.  |
| 419H   | 1049 | IA32_MC6_STATUS            | Package |  |
| 41AH   | 1050 | IA32_MC6_ADDR              | Package |  |
| 41BH   | 1051 | IA32_MC6_MISC              | Package |  |
| 41CH   | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC7 reports MC errors from the home agent HA 0.  |
| 41DH   | 1053 | IA32_MC7_STATUS            | Package |  |
| 41EH   | 1054 | IA32_MC7_ADDR              | Package |  |
| 41FH   | 1055 | IA32_MC7_MISC              | Package |  |
| 424H   | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 10 report MC errors from each channel of the integrated memory controllers.                                      |
| 425H   | 1061 | IA32_MC9_STATUS            | Package |  |
| 426H   | 1062 | IA32_MC9_ADDR              | Package |  |
| 427H   | 1063 | IA32_MC9_MISC              | Package |  |
| 428H   | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 10 report MC errors from each channel of the integrated memory controllers.                                      |
| 429H   | 1065 | IA32_MC10_STATUS           | Package |  |
| 42AH   | 1066 | IA32_MC10_ADDR             | Package |  |
| 42BH   | 1067 | IA32_MC10_MISC             | Package |  |
| 444H   | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC17 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo0, CBo3, CBo6, CBo9, CBo12, CBo15.  |
| 445H   | 1093 | IA32_MC17_STATUS           | Package |  |
| 446H   | 1094 | IA32_MC17_ADDR             | Package |  |
| 447H   | 1095 | IA32_MC17_MISC             | Package |  |
| 448H   | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC18 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo1, CBo4, CBo7, CBo10, CBo13, CBo16. |
| 449H   | 1097 | IA32_MC18_STATUS           | Package |  |
| 44AH   | 1098 | IA32_MC18_ADDR             | Package |  |
| 44BH   | 1099 | IA32_MC18_MISC             | Package |  |
| 44CH   | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC19 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo2, CBo5, CBo8, CBo11, CBo14, CBo17. |
| 44DH   | 1101 | IA32_MC19_STATUS           | Package |  |
| 44EH   | 1102 | IA32_MC19_ADDR             | Package |  |
| 44FH   | 1103 | IA32_MC19_MISC             | Package |  |
| See Table 2-19, Table 2-28, Table 2-33, and Table 2-35 for other MSR definitions applicable to processors with CPUID signature 06_56H. |      |                            |         |  |

**NOTES:**

1. An override configuration lower than the factory-set configuration is always supported. An override configuration higher than the factory-set configuration is dependent on features specific to the processor and the platform.

**2.15.2 Additional MSRs Supported in Intel® Xeon® Processors E5 v4 and E7 v4 Families**

The MSRs listed in Table 2-36 are available to Intel® Xeon® Processor E5 v4 and E7 v4 Families (CPUID DisplayFamily\_DisplayModel = 06\_4FH). The Intel® Xeon® processor E5 v4 family is based on the Broadwell micro-



architecture and supports the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-28, Table 2-33, Table 2-35, and Table 2-37.

**Table 2-37. Additional MSRs Supported by Intel® Xeon® Processors with DisplayFamily\_DisplayModel 06\_4FH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 1ACH             | 428  | MSR_TURBO_RATIO_LIMIT3     | Package | Config Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0.<br>RW if MSR_PLATFORM_INFO.[28] = 1.  |
|                  |      | 62:0                       | Package | Reserved  |
|                  |      | 63                         | Package | Semaphore for Turbo Ratio Limit Configuration<br>If 1, the processor uses override configuration <sup>1</sup> specified in MSR_TURBO_RATIO_LIMIT, MSR_TURBO_RATIO_LIMIT1 and MSR_TURBO_RATIO_LIMIT2.<br>If 0, the processor uses factory-set configuration (Default). |
| 285H             | 645  | IA32_MC5_CTL2              | Package | See Table 2-2.  |
| 286H             | 646  | IA32_MC6_CTL2              | Package | See Table 2-2.  |
| 287H             | 647  | IA32_MC7_CTL2              | Package | See Table 2-2.  |
| 288H             | 648  | IA32_MC8_CTL2              | Package | See Table 2-2.  |
| 289H             | 649  | IA32_MC9_CTL2              | Package | See Table 2-2.  |
| 28AH             | 650  | IA32_MC10_CTL2             | Package | See Table 2-2.  |
| 28BH             | 651  | IA32_MC11_CTL2             | Package | See Table 2-2.  |
| 28CH             | 652  | IA32_MC12_CTL2             | Package | See Table 2-2.  |
| 28DH             | 653  | IA32_MC13_CTL2             | Package | See Table 2-2.  |
| 28EH             | 654  | IA32_MC14_CTL2             | Package | See Table 2-2.  |
| 28FH             | 655  | IA32_MC15_CTL2             | Package | See Table 2-2.  |
| 290H             | 656  | IA32_MC16_CTL2             | Package | See Table 2-2.  |
| 291H             | 657  | IA32_MC17_CTL2             | Package | See Table 2-2.  |
| 292H             | 658  | IA32_MC18_CTL2             | Package | See Table 2-2.  |
| 293H             | 659  | IA32_MC19_CTL2             | Package | See Table 2-2.  |
| 294H             | 660  | IA32_MC20_CTL2             | Package | See Table 2-2.  |
| 295H             | 661  | IA32_MC21_CTL2             | Package | See Table 2-2.  |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC5 reports MC errors from the Intel QPI 0 module.  |
| 415H             | 1045 | IA32_MC5_STATUS            | Package |   |
| 416H             | 1046 | IA32_MC5_ADDR              | Package |   |
| 417H             | 1047 | IA32_MC5_MISC              | Package |   |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC6 reports MC errors from the integrated I/O module.   |
| 419H             | 1049 | IA32_MC6_STATUS            | Package |   |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package |   |
| 41BH             | 1051 | IA32_MC6_MISC              | Package |   |

**Table 2-37. Additional MSRs Supported by Intel® Xeon® Processors with DisplayFamily\_DisplayModel 06\_4FH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC7 reports MC errors from the home agent HA 0.   |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package |   |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package |   |
| 41FH             | 1055 | IA32_MC7_MISC              | Package |   |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC8 reports MC errors from the home agent HA 1.   |
| 421H             | 1057 | IA32_MC8_STATUS            | Package |   |
| 422H             | 1058 | IA32_MC8_ADDR              | Package |   |
| 423H             | 1059 | IA32_MC8_MISC              | Package |   |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 425H             | 1061 | IA32_MC9_STATUS            | Package |   |
| 426H             | 1062 | IA32_MC9_ADDR              | Package |   |
| 427H             | 1063 | IA32_MC9_MISC              | Package |   |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 429H             | 1065 | IA32_MC10_STATUS           | Package |   |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package |   |
| 42BH             | 1067 | IA32_MC10_MISC             | Package |   |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package |   |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package |   |
| 42FH             | 1071 | IA32_MC11_MISC             | Package |   |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 431H             | 1073 | IA32_MC12_STATUS           | Package |   |
| 432H             | 1074 | IA32_MC12_ADDR             | Package |   |
| 433H             | 1075 | IA32_MC12_MISC             | Package |   |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 435H             | 1077 | IA32_MC13_STATUS           | Package |   |
| 436H             | 1078 | IA32_MC13_ADDR             | Package |   |
| 437H             | 1079 | IA32_MC13_MISC             | Package |   |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 439H             | 1081 | IA32_MC14_STATUS           | Package |   |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package |   |
| 43BH             | 1083 | IA32_MC14_MISC             | Package |   |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers. |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package |   |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package |   |
| 43FH             | 1087 | IA32_MC15_MISC             | Package |   |

**Table 2-37. Additional MSRs Supported by Intel® Xeon® Processors with DisplayFamily\_DisplayModel 06\_4FH**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 through MC 16 report MC errors from each channel of the integrated memory controllers.                                      |
| 441H             | 1089 | IA32_MC16_STATUS           | Package |  |
| 442H             | 1090 | IA32_MC16_ADDR             | Package |  |
| 443H             | 1091 | IA32_MC16_MISC             | Package |  |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC17 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo0, CBo3, CBo6, CBo9, CBo12, CBo15.  |
| 445H             | 1093 | IA32_MC17_STATUS           | Package |  |
| 446H             | 1094 | IA32_MC17_ADDR             | Package |  |
| 447H             | 1095 | IA32_MC17_MISC             | Package |  |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC18 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo1, CBo4, CBo7, CBo10, CBo13, CBo16. |
| 449H             | 1097 | IA32_MC18_STATUS           | Package |  |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package |  |
| 44BH             | 1099 | IA32_MC18_MISC             | Package |  |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC19 reports MC errors from the following pair of CBo/L3 Slices (if the pair is present): CBo2, CBo5, CBo8, CBo11, CBo14, CBo17. |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package |  |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package |  |
| 44FH             | 1103 | IA32_MC19_MISC             | Package |  |
| 450H             | 1104 | IA32_MC20_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC20 reports MC errors from the Intel QPI 1 module.  |
| 451H             | 1105 | IA32_MC20_STATUS           | Package |  |
| 452H             | 1106 | IA32_MC20_ADDR             | Package |  |
| 453H             | 1107 | IA32_MC20_MISC             | Package |  |
| 454H             | 1108 | IA32_MC21_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC21 reports MC errors from the Intel QPI 2 module.  |
| 455H             | 1109 | IA32_MC21_STATUS           | Package |  |
| 456H             | 1110 | IA32_MC21_ADDR             | Package |  |
| 457H             | 1111 | IA32_MC21_MISC             | Package |  |
| C81H             | 3201 | IA32_L3_QOS_CFG            | Package | Cache Allocation Technology Configuration (R/W)  |
|                  |      | 0                          |         | CAT Enable. Set 1 to enable Cache Allocation Technology.   |
|                  |      | 63:1                       |         | Reserved   |

See Table 2-19, Table 2-20, Table 2-28, and Table 2-29 for other MSR definitions applicable to processors with CPUID signature 06\_45H.

**NOTES:**

1. An override configuration lower than the factory-set configuration is always supported. An override configuration higher than the factory-set configuration is dependent on features specific to the processor and the platform.

## 2.16 MSRS IN THE 6TH GENERATION, 7TH GENERATION AND 8TH GENERATION INTEL® CORE™ PROCESSORS, INTEL® XEON® PROCESSOR SCALABLE FAMILY, AND FUTURE INTEL® CORE™ PROCESSORS

6th generation Intel® Core™ processors and the Intel® Xeon® Processor Scalable Family are based on the Skylake microarchitecture and have CPUID DisplayFamily\_DisplayModel signatures of 06\_4EH, 06\_5EH, and 06\_55H. 7th generation Intel® Core™ processors are based on the Kaby Lake microarchitecture and 8th generation Intel® Core™ processors are based on the Coffee Lake microarchitecture; these processors have CPUID DisplayFamily\_DisplayModel signatures of 06\_8EH and 06\_9EH. Future Intel® Core™ processors are based on Cannon Lake microarchitecture and have a CPUID DisplayFamily\_DisplayModel signature of 06\_66H. These processors support the MSR interfaces listed in Table 2-19, Table 2-20, Table 2-24, Table 2-28, Table 2-34, Table 2-38, and Table 2-39. For an MSR listed in Table 2-38 that also appears in the model-specific tables of prior generations, Table 2-38 supercede prior generation tables.

The notation of “Platform” in the Scope column (with respect to MSR\_PLATFORM\_ENERGY\_COUNTER and MSR\_PLATFORM\_POWER\_LIMIT) is limited to the power-delivery domain and the specifics of the power delivery integration may vary by platform vendor’s implementation.

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2. |
| FEH              | 254 | IA32_MTRRCAP               | Thread | MTRR Capability (RO, Architectural)<br>See Table 2-2           |
| 19CH             | 412 | IA32_THERM_STATUS          | Core   | Thermal Monitor Status (R/W)<br>See Table 2-2.                 |
|                  |     | 0                          |        | Thermal Status (RO)<br>See Table 2-2.                          |
|                  |     | 1                          |        | Thermal Status Log (R/WCO)<br>See Table 2-2.                   |
|                  |     | 2                          |        | PROTCHOT # or FORCEPR# Status (RO)<br>See Table 2-2.           |
|                  |     | 3                          |        | PROTCHOT # or FORCEPR# Log (R/WCO)<br>See Table 2-2.           |
|                  |     | 4                          |        | Critical Temperature Status (RO)<br>See Table 2-2.             |
|                  |     | 5                          |        | Critical Temperature Status Log (R/WCO)<br>See Table 2-2.      |
|                  |     | 6                          |        | Thermal threshold #1 Status (RO)<br>See Table 2-2.             |
|                  |     | 7                          |        | Thermal threshold #1 Log (R/WCO)<br>See Table 2-2.             |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 8                          |         | Thermal Threshold #2 Status (RO)<br>See Table 2-2.   |
|                  |     | 9                          |         | Thermal Threshold #2 Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 10                         |         | Power Limitation Status (RO)<br>See Table 2-2.   |
|                  |     | 11                         |         | Power Limitation Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 12                         |         | Current Limit Status (RO)<br>See Table 2-2.  |
|                  |     | 13                         |         | Current Limit Log (R/WCO)<br>See Table 2-2.  |
|                  |     | 14                         |         | Cross Domain Limit Status (RO)<br>See Table 2-2.   |
|                  |     | 15                         |         | Cross Domain Limit Log (R/WCO)<br>See Table 2-2.   |
|                  |     | 22:16                      |         | Digital Readout (RO)<br>See Table 2-2.   |
|                  |     | 26:23                      |         | Reserved   |
|                  |     | 30:27                      |         | Resolution in Degrees Celsius (RO)<br>See Table 2-2.   |
|                  |     | 31                         |         | Reading Valid (RO)<br>See Table 2-2.   |
|                  |     | 63:32                      |         | Reserved   |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode<br>RO if MSR_PLATFORM_INFO.[28] = 0,<br>RW if MSR_PLATFORM_INFO.[28] = 1 |
|                  |     | 7:0                        | Package | Maximum Ratio Limit for 1C<br>Maximum turbo ratio limit of 1 core active.                                  |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for 2C<br>Maximum turbo ratio limit of 2 core active.                                  |
|                  |     | 23:16                      | Package | Maximum Ratio Limit for 3C<br>Maximum turbo ratio limit of 3 core active.                                  |
|                  |     | 31:24                      | Package | Maximum Ratio Limit for 4C<br>Maximum turbo ratio limit of 4 core active.                                  |
|                  |     | 63:32                      |         | Reserved   |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS         | Thread  | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-4) that points to the MSR containing the most recent branch record.   |
| 1FCH             | 508 | MSR_POWER_CTL              | Core    | Power Control Register<br>See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 0                          |         | Reserved  |
|                  |     | 1                          | Package | C1E Enable (R/W)<br>When set to '1', will enable the CPU to switch to the Minimum Enhanced Intel SpeedStep Technology operating point when all execution cores enter MWAIT (C1).  |
|                  |     | 18:2                       |         | Reserved  |
|                  |     | 19                         |         | Disable Race to Halt Optimization (R/W)<br>Setting this bit disables the Race to Halt optimization and avoids this optimization limitation to execute below the most efficient frequency ratio. Default value is 0 for processors that support Race to Halt optimization. Default value is 1 for processors that do not support Race to Halt optimization.  |
|                  |     | 20                         |         | Disable Energy Efficiency Optimization (R/W)<br>Setting this bit disables the P-States energy efficiency optimization. Default value is 0. Disable/enable the energy efficiency optimization in P-State legacy mode (when IA32_PM_ENABLE[HWP_ENABLE] = 0), has an effect only in the turbo range or into PERF_MIN_CTL value if it is not zero set. In HWP mode (IA32_PM_ENABLE[HWP_ENABLE] == 1), has an effect between the OS desired or OS maximize to the OS minimize performance setting. |
| 300H             | 768 | MSR_SGXOWNEREPOCH0         | Package | Lower 64 Bit CR_SGXOWNEREPOCH (W)<br>Writes do not update CR_SGXOWNEREPOCH if CPUID.(EAX=12H, ECX=0):EAX.SGX1 is 1 on any thread in the package.  |
|                  |     | 63:0                       |         | Lower 64 bits of an 128-bit external entropy value for key derivation of an enclave.  |
| 301H             | 768 | MSR_SGXOWNEREPOCH1         | Package | Upper 64 Bit CR_SGXOWNEREPOCH (W)<br>Writes do not update CR_SGXOWNEREPOCH if CPUID.(EAX=12H, ECX=0):EAX.SGX1 is 1 on any thread in the package.  |
|                  |     | 63:0                       |         | Upper 64 bits of an 128-bit external entropy value for key derivation of an enclave.  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |        | Register Name / Bit Fields    | Scope  | Bit Description  |
|------------------|--------|-------------------------------|--------|--|
| Hex              | Dec    |                               |        |  |
| 38EH             | 910    | IA32_PERF_GLOBAL_STATUS       |        | See Table 2-2. See Section 18.2.4, “Architectural Performance Monitoring Version 4.” |
|                  |        | 0                             | Thread | Ovf_PMC0   |
|                  |        | 1                             | Thread | Ovf_PMC1   |
|                  |        | 2                             | Thread | Ovf_PMC2   |
|                  |        | 3                             | Thread | Ovf_PMC3   |
|                  |        | 4                             | Thread | Ovf_PMC4 (if CPUID.0AH:EAX[15:8] > 4)  |
|                  |        | 5                             | Thread | Ovf_PMC5 (if CPUID.0AH:EAX[15:8] > 5)  |
|                  |        | 6                             | Thread | Ovf_PMC6 (if CPUID.0AH:EAX[15:8] > 6)  |
|                  |        | 7                             | Thread | Ovf_PMC7 (if CPUID.0AH:EAX[15:8] > 7)  |
|                  |        | 31:8                          |        | Reserved   |
|                  |        | 32                            | Thread | Ovf_FixedCtr0  |
|                  |        | 33                            | Thread | Ovf_FixedCtr1  |
|                  |        | 34                            | Thread | Ovf_FixedCtr2  |
|                  |        | 54:35                         |        | Reserved   |
|                  |        | 55                            | Thread | Trace_ToPA_PMI   |
|                  |        | 57:56                         |        | Reserved   |
|                  |        | 58                            | Thread | LBR_Frz  |
|                  |        | 59                            | Thread | CTR_Frz  |
|                  |        | 60                            | Thread | ASCI   |
|                  |        | 61                            | Thread | Ovf_Uncore   |
| 62               | Thread | Ovf_BufDSSAVE                 |        |  |
| 63               | Thread | CondChgd                      |        |  |
| 390H             | 912    | IA32_PERF_GLOBAL_STATUS_RESET |        | See Table 2-2. See Section 18.2.4, “Architectural Performance Monitoring Version 4.” |
|                  |        | 0                             | Thread | Set 1 to clear Ovf_PMC0.   |
|                  |        | 1                             | Thread | Set 1 to clear Ovf_PMC1.   |
|                  |        | 2                             | Thread | Set 1 to clear Ovf_PMC2.   |
|                  |        | 3                             | Thread | Set 1 to clear Ovf_PMC3.   |
|                  |        | 4                             | Thread | Set 1 to clear Ovf_PMC4 (if CPUID.0AH:EAX[15:8] > 4).                                |
|                  |        | 5                             | Thread | Set 1 to clear Ovf_PMC5 (if CPUID.0AH:EAX[15:8] > 5).                                |
|                  |        | 6                             | Thread | Set 1 to clear Ovf_PMC6 (if CPUID.0AH:EAX[15:8] > 6).                                |
|                  |        | 7                             | Thread | Set 1 to clear Ovf_PMC7 (if CPUID.0AH:EAX[15:8] > 7).                                |
|                  |        | 31:8                          |        | Reserved   |
| 32               | Thread | Set 1 to clear Ovf_FixedCtr0. |        |  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields  | Scope  | Bit Description  |
|------------------|-----|-----------------------------|--------|--|
| Hex              | Dec |                             |        |  |
|                  |     | 33                          | Thread | Set 1 to clear Ovf_FixedCtr1.  |
|                  |     | 34                          | Thread | Set 1 to clear Ovf_FixedCtr2.  |
|                  |     | 54:35                       |        | Reserved   |
|                  |     | 55                          | Thread | Set 1 to clear Trace_ToPA_PMI.   |
|                  |     | 57:56                       |        | Reserved   |
|                  |     | 58                          | Thread | Set 1 to clear LBR_Frz.  |
|                  |     | 59                          | Thread | Set 1 to clear CTR_Frz.  |
|                  |     | 60                          | Thread | Set 1 to clear ASCI.   |
|                  |     | 61                          | Thread | Set 1 to clear Ovf_Uncore.   |
|                  |     | 62                          | Thread | Set 1 to clear Ovf_BufDSSAVE.  |
|                  |     | 63                          | Thread | Set 1 to clear CondChgd.   |
| 391H             | 913 | IA32_PERF_GLOBAL_STATUS_SET |        | See Table 2-2. See Section 18.2.4, "Architectural Performance Monitoring Version 4." |
|                  |     | 0                           | Thread | Set 1 to cause Ovf_PMC0 = 1.   |
|                  |     | 1                           | Thread | Set 1 to cause Ovf_PMC1 = 1.   |
|                  |     | 2                           | Thread | Set 1 to cause Ovf_PMC2 = 1.   |
|                  |     | 3                           | Thread | Set 1 to cause Ovf_PMC3 = 1.   |
|                  |     | 4                           | Thread | Set 1 to cause Ovf_PMC4=1 (if CPUID.0AH:EAX[15:8] > 4).                              |
|                  |     | 5                           | Thread | Set 1 to cause Ovf_PMC5=1 (if CPUID.0AH:EAX[15:8] > 5).                              |
|                  |     | 6                           | Thread | Set 1 to cause Ovf_PMC6=1 (if CPUID.0AH:EAX[15:8] > 6).                              |
|                  |     | 7                           | Thread | Set 1 to cause Ovf_PMC7=1 (if CPUID.0AH:EAX[15:8] > 7).                              |
|                  |     | 31:8                        |        | Reserved   |
|                  |     | 32                          | Thread | Set 1 to cause Ovf_FixedCtr0 = 1.  |
|                  |     | 33                          | Thread | Set 1 to cause Ovf_FixedCtr1 = 1.  |
|                  |     | 34                          | Thread | Set 1 to cause Ovf_FixedCtr2 = 1.  |
|                  |     | 54:35                       |        | Reserved   |
|                  |     | 55                          | Thread | Set 1 to cause Trace_ToPA_PMI = 1.   |
|                  |     | 57:56                       |        | Reserved   |
|                  |     | 58                          | Thread | Set 1 to cause LBR_Frz = 1.  |
|                  |     | 59                          | Thread | Set 1 to cause CTR_Frz = 1.  |
|                  |     | 60                          | Thread | Set 1 to cause ASCI = 1.   |
|                  |     | 61                          | Thread | Set 1 to cause Ovf_Uncore.   |



**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
|                  |      | 62                         | Thread | Set 1 to cause Ovf_BufDSSAVE.  |
|                  |      | 63                         |        | Reserved   |
| 392H             | 913  | IA32_PERF_GLOBAL_INUSE     |        | See Table 2-2.   |
| 3F7H             | 1015 | MSR_PEBS_FRONTEND          | Thread | FrontEnd Precise Event Condition Select (R/W)  |
|                  |      | 2:0                        |        | Event Code Select  |
|                  |      | 3                          |        | Reserved   |
|                  |      | 4                          |        | Event Code Select High   |
|                  |      | 7:5                        |        | Reserved   |
|                  |      | 19:8                       |        | IDQ_Bubble_Length Specifier  |
|                  |      | 22:20                      |        | IDQ_Bubble_Width Specifier   |
|                  |      | 63:23                      |        | Reserved   |
| 500H             | 1280 | IA32_SGX_SVN_STATUS        | Thread | Status and SVN Threshold of SGX Support for ACM (RO)   |
|                  |      | 0                          |        | Lock<br>See Section 41.11.3, "Interactions with Authenticated Code Modules (ACMs)".          |
|                  |      | 15:1                       |        | Reserved   |
|                  |      | 23:16                      |        | SGX_SVN_SINIT<br>See Section 41.11.3, "Interactions with Authenticated Code Modules (ACMs)". |
|                  |      | 63:24                      |        | Reserved   |
| 560H             | 1376 | IA32_RTIT_OUTPUT_BASE      | Thread | Trace Output Base Register (R/W)<br>See Table 2-2.   |
| 561H             | 1377 | IA32_RTIT_OUTPUT_MASK_PTRS | Thread | Trace Output Mask Pointers Register (R/W)<br>See Table 2-2.                                  |
| 570H             | 1392 | IA32_RTIT_CTL              | Thread | Trace Control Register (R/W)   |
|                  |      | 0                          |        | TraceEn  |
|                  |      | 1                          |        | CYCEn  |
|                  |      | 2                          |        | OS   |
|                  |      | 3                          |        | User   |
|                  |      | 6:4                        |        | Reserved, must be zero.  |
|                  |      | 7                          |        | CR3 filter   |
|                  |      | 8                          |        | ToPA<br>Writing 0 will #GP if also setting TraceEn.  |
|                  |      | 9                          |        | MTCEn  |
|                  |      | 10                         |        | TSCEn  |
|                  |      | 11                         |        | DisRETC  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 12                         |         | Reserved, must be zero.  |
|                  |      | 13                         |         | BranchEn   |
|                  |      | 17:14                      |         | MTCFreq  |
|                  |      | 18                         |         | Reserved, must be zero.  |
|                  |      | 22:19                      |         | CYCThresh  |
|                  |      | 23                         |         | Reserved, must be zero.  |
|                  |      | 27:24                      |         | PSBFreq  |
|                  |      | 31:28                      |         | Reserved, must be zero.  |
|                  |      | 35:32                      |         | ADDR0_CFG  |
|                  |      | 39:36                      |         | ADDR1_CFG  |
|                  |      | 63:40                      |         | Reserved, must be zero.  |
| 571H             | 1393 | IA32_RTIT_STATUS           | Thread  | Tracing Status Register (R/W)  |
|                  |      | 0                          |         | FilterEn, writes ignored.  |
|                  |      | 1                          |         | ContexEn, writes ignored.  |
|                  |      | 2                          |         | TriggerEn, writes ignored.   |
|                  |      | 3                          |         | Reserved   |
|                  |      | 4                          |         | Error (R/W)  |
|                  |      | 5                          |         | Stopped  |
|                  |      | 31:6                       |         | Reserved, must be zero.  |
|                  |      | 48:32                      |         | PacketByteCnt  |
|                  |      | 63:49                      |         | Reserved, must be zero.  |
| 572H             | 1394 | IA32_RTIT_CR3_MATCH        | Thread  | Trace Filter CR3 Match Register (R/W)                                  |
|                  |      | 4:0                        |         | Reserved   |
|                  |      | 63:5                       |         | CR3[63:5] value to match   |
| 580H             | 1408 | IA32_RTIT_ADDR0_A          | Thread  | Region 0 Start Address (R/W)   |
|                  |      | 63:0                       |         | See Table 2-2.   |
| 581H             | 1409 | IA32_RTIT_ADDR0_B          | Thread  | Region 0 End Address (R/W)   |
|                  |      | 63:0                       |         | See Table 2-2.   |
| 582H             | 1410 | IA32_RTIT_ADDR1_A          | Thread  | Region 1 Start Address (R/W)   |
|                  |      | 63:0                       |         | See Table 2-2.   |
| 583H             | 1411 | IA32_RTIT_ADDR1_B          | Thread  | Region 1 End Address (R/W)   |
|                  |      | 63:0                       |         | See Table 2-2.   |
| 639H             | 1593 | MSR_PP0_ENERGY_STATUS      | Package | PP0 Energy Status (R/O)<br>See Section 14.9.4, "PP0/PP1 RAPL Domains." |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields  | Scope     | Bit Description  |
|------------------|------|-----------------------------|-----------|--|
| Hex              | Dec  |                             |           |  |
| 64DH             | 1613 | MSR_PLATFORM_ENERGY_COUNTER | Platform* | Platform Energy Counter (R/O)<br>This MSR is valid only if both platform vendor hardware implementation and BIOS enablement support it. This MSR will read 0 if not valid.   |
|                  |      | 31:0                        |           | Total energy consumed by all devices in the platform that receive power from integrated power delivery mechanism, included platform devices are processor cores, SOC, memory, add-on or peripheral devices that get powered directly from the platform power delivery means. The energy units are specified in the MSR_RAPL_POWER_UNIT.Energy_Status_Unit. |
|                  |      | 63:32                       |           | Reserved   |
| 64EH             | 1614 | MSR_PPERF                   | Thread    | Productive Performance Count (R/O)   |
|                  |      | 63:0                        |           | Hardware's view of workload scalability. See Section 14.4.5.1.   |
| 64FH             | 1615 | MSR_CORE_PERF_LIMIT_REASONS | Package   | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.)  |
|                  |      | 0                           |           | PROCHOT Status (R0)<br>When set, frequency is reduced below the operating system request due to assertion of external PROCHOT.   |
|                  |      | 1                           |           | Thermal Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal event.   |
|                  |      | 3:2                         |           | Reserved   |
|                  |      | 4                           |           | Residency State Regulation Status (R0)<br>When set, frequency is reduced below the operating system request due to residency state regulation limit.   |
|                  |      | 5                           |           | Running Average Thermal Limit Status (R0)<br>When set, frequency is reduced below the operating system request due to Running Average Thermal Limit (RATL).  |
|                  |      | 6                           |           | VR Therm Alert Status (R0)<br>When set, frequency is reduced below the operating system request due to a thermal alert from a processor Voltage Regulator (VR).  |
|                  |      | 7                           |           | VR Therm Design Current Status (R0)<br>When set, frequency is reduced below the operating system request due to VR thermal design current limit.   |
|                  |      | 8                           |           | Other Status (R0)<br>When set, frequency is reduced below the operating system request due to electrical or other constraints.   |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 9                          |       | Reserved  |
|                  |     | 10                         |       | Package/Platform-Level Power Limiting PL1 Status (RO)<br>When set, frequency is reduced below the operating system request due to package/platform-level power limiting PL1.  |
|                  |     | 11                         |       | Package/Platform-Level PL2 Power Limiting Status (RO)<br>When set, frequency is reduced below the operating system request due to package/platform-level power limiting PL2/PL3.  |
|                  |     | 12                         |       | Max Turbo Limit Status (RO)<br>When set, frequency is reduced below the operating system request due to multi-core turbo limits.  |
|                  |     | 13                         |       | Turbo Transition Attenuation Status (RO)<br>When set, frequency is reduced below the operating system request due to Turbo transition attenuation. This prevents performance degradation due to frequent operating ratio changes. |
|                  |     | 15:14                      |       | Reserved  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |     | 19:18                      |       | Reserved.   |
|                  |     | 20                         |       | Residency State Regulation Log<br>When set, indicates that the Residency State Regulation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.        |
|                  |     | 21                         |       | Running Average Thermal Limit Log<br>When set, indicates that the RATL Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                           |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 22                         |         | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 23                         |         | VR Thermal Design Current Log<br>When set, indicates that the VR TDC Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 24                         |         | Other Log<br>When set, indicates that the Other Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 25                         |         | Reserved  |
|                  |      | 26                         |         | Package/Platform-Level PL1 Power Limiting Log<br>When set, indicates that the Package or Platform Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.     |
|                  |      | 27                         |         | Package/Platform-Level PL2 Power Limiting Log<br>When set, indicates that the Package or Platform Level PL2/PL3 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 28                         |         | Max Turbo Limit Log<br>When set, indicates that the Max Turbo Limit Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |      | 29                         |         | Turbo Transition Attenuation Log<br>When set, indicates that the Turbo Transition Attenuation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                                  |
|                  |      | 63:30                      |         | Reserved  |
| 652H             | 1618 | MSR_PKG_HDC_CONFIG         | Package | HDC Configuration (R/W)   |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields    | Scope     | Bit Description   |
|------------------|------|-------------------------------|-----------|---|
| Hex              | Dec  |                               |           |   |
|                  |      | 2:0                           |           | PKG_Cx_Monitor<br>Configures Package Cx state threshold for MSR_PKG_HDC_DEEP_RESIDENCY.   |
|                  |      | 63:3                          |           | Reserved  |
| 653H             | 1619 | MSR_CORE_HDC_RESIDENCY        | Core      | Core HDC Idle Residency (R/O)   |
|                  |      | 63:0                          |           | Core_Cx_Duty_Cycle_Cnt  |
| 655H             | 1621 | MSR_PKG_HDC_SHALLOW_RESIDENCY | Package   | Accumulate the cycles the package was in C2 state and at least one logical processor was in forced idle (R/O)   |
|                  |      | 63:0                          |           | Pkg_C2_Duty_Cycle_Cnt   |
| 656H             | 1622 | MSR_PKG_HDC_DEEP_RESIDENCY    | Package   | Package Cx HDC Idle Residency (R/O)   |
|                  |      | 63:0                          |           | Pkg_Cx_Duty_Cycle_Cnt   |
| 658H             | 1624 | MSR_WEIGHTED_CORE_CO          | Package   | Core-count Weighted C0 Residency (R/O)  |
|                  |      | 63:0                          |           | Increment at the same rate as the TSC. The increment each cycle is weighted by the number of processor cores in the package that reside in C0. If N cores are simultaneously in C0, then each cycle the counter increments by N.  |
| 659H             | 1625 | MSR_ANY_CORE_CO               | Package   | Any Core C0 Residency (R/O)   |
|                  |      | 63:0                          |           | Increment at the same rate as the TSC. The increment each cycle is one if any processor core in the package is in C0.   |
| 65AH             | 1626 | MSR_ANY_GFXE_CO               | Package   | Any Graphics Engine C0 Residency (R/O)  |
|                  |      | 63:0                          |           | Increment at the same rate as the TSC. The increment each cycle is one if any processor graphic device's compute engines are in C0.   |
| 65BH             | 1627 | MSR_CORE_GFXE_OVERLAP_CO      | Package   | Core and Graphics Engine Overlapped C0 Residency (R/O)  |
|                  |      | 63:0                          |           | Increment at the same rate as the TSC. The increment each cycle is one if at least one compute engine of the processor graphics is in C0 and at least one processor core in the package is also in C0.  |
| 65CH             | 1628 | MSR_PLATFORM_POWER_LIMIT      | Platform* | Platform Power Limit Control (R/W-L)<br>Allows platform BIOS to limit power consumption of the platform devices to the specified values. The Long Duration power consumption is specified via Platform_Power_Limit_1 and Platform_Power_Limit_1_Time. The Short Duration power consumption limit is specified via the Platform_Power_Limit_2 with duration chosen by the processor.<br><br>The processor implements an exponential-weighted algorithm in the placement of the time windows. |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 14:0                       |       | Platform Power Limit #1<br>Average Power limit value which the platform must not exceed over a time window as specified by Power_Limit_1_TIME field.<br>The default value is the Thermal Design Power (TDP) and varies with product skus. The unit is specified in MSR_RAPLPOWER_UNIT.   |
|                  |     | 15                         |       | Enable Platform Power Limit #1<br>When set, enables the processor to apply control policy such that the platform power does not exceed Platform Power limit #1 over the time window specified by Power Limit #1 Time Window.   |
|                  |     | 16                         |       | Platform Clamping Limitation #1<br>When set, allows the processor to go below the OS requested P states in order to maintain the power below specified Platform Power Limit #1 value.<br>This bit is writeable only when CPUID (EAX=6):EAX[4] is set.  |
|                  |     | 23:17                      |       | Time Window for Platform Power Limit #1<br>Specifies the duration of the time window over which Platform Power Limit 1 value should be maintained for sustained long duration. This field is made up of two numbers from the following equation:<br>Time Window = (float) ((1+(X/4))*(2^Y)), where:<br>X = POWER_LIMIT_1_TIME[23:22]<br>Y = POWER_LIMIT_1_TIME[21:17]<br>The maximum allowed value in this field is defined in MSR_PKG_POWER_INFO[PKG_MAX_WIN].<br>The default value is 0DH, The unit is specified in MSR_RAPLPOWER_UNIT[Time Unit]. |
|                  |     | 31:24                      |       | Reserved   |
|                  |     | 46:32                      |       | Platform Power Limit #2<br>Average Power limit value which the platform must not exceed over the Short Duration time window chosen by the processor.<br>The recommended default value is 1.25 times the Long Duration Power Limit (i.e., Platform Power Limit # 1).  |
|                  |     | 47                         |       | Enable Platform Power Limit #2<br>When set, enables the processor to apply control policy such that the platform power does not exceed Platform Power limit #2 over the Short Duration time window.  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
|                  |      | 48                         |        | Platform Clamping Limitation #2<br>When set, allows the processor to go below the OS requested P states in order to maintain the power below specified Platform Power Limit #2 value.  |
|                  |      | 62:49                      |        | Reserved   |
|                  |      | 63                         |        | Lock. Setting this bit will lock all other bits of this MSR until system RESET.  |
| 690H             | 1680 | MSR_LASTBRANCH_16_FROM_IP  | Thread | Last Branch Record 16 From IP (R/W)<br>One of 32 triplets of last branch record registers on the last branch record stack. This part of the stack contains pointers to the source instruction. See also: <ul style="list-style-type: none"> <li>Last Branch Record Stack TOS at 1C9H.</li> <li>Section 17.12.</li> </ul> |
| 691H             | 1681 | MSR_LASTBRANCH_17_FROM_IP  | Thread | Last Branch Record 17 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 692H             | 1682 | MSR_LASTBRANCH_18_FROM_IP  | Thread | Last Branch Record 18 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 693H             | 1683 | MSR_LASTBRANCH_19_FROM_IP  | Thread | Last Branch Record 19 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 694H             | 1684 | MSR_LASTBRANCH_20_FROM_IP  | Thread | Last Branch Record 20 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 695H             | 1685 | MSR_LASTBRANCH_21_FROM_IP  | Thread | Last Branch Record 21 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 696H             | 1686 | MSR_LASTBRANCH_22_FROM_IP  | Thread | Last Branch Record 22 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 697H             | 1687 | MSR_LASTBRANCH_23_FROM_IP  | Thread | Last Branch Record 23 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 698H             | 1688 | MSR_LASTBRANCH_24_FROM_IP  | Thread | Last Branch Record 24 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 699H             | 1689 | MSR_LASTBRANCH_25_FROM_IP  | Thread | Last Branch Record 25 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69AH             | 1690 | MSR_LASTBRANCH_26_FROM_IP  | Thread | Last Branch Record 26 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69BH             | 1691 | MSR_LASTBRANCH_27_FROM_IP  | Thread | Last Branch Record 27 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |
| 69CH             | 1692 | MSR_LASTBRANCH_28_FROM_IP  | Thread | Last Branch Record 28 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.  |



**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields      | Scope   | Bit Description   |
|------------------|------|---------------------------------|---------|---|
| Hex              | Dec  |                                 |         |   |
| 69DH             | 1693 | MSR_LASTBRANCH_29_FROM_IP       | Thread  | Last Branch Record 29 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 69EH             | 1694 | MSR_LASTBRANCH_30_FROM_IP       | Thread  | Last Branch Record 30 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 69FH             | 1695 | MSR_LASTBRANCH_31_FROM_IP       | Thread  | Last Branch Record 31 From IP (R/W)<br>See description of MSR_LASTBRANCH_0_FROM_IP.   |
| 6B0H             | 1712 | MSR_GRAPHICS_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in the Processor Graphics (R/W)<br>(Frequency refers to processor graphics frequency.)                        |
|                  |      | 0                               |         | PROCHOT Status (RO)<br>When set, frequency is reduced due to assertion of external PROCHOT.   |
|                  |      | 1                               |         | Thermal Status (RO)<br>When set, frequency is reduced due to a thermal event.   |
|                  |      | 4:2                             |         | Reserved.   |
|                  |      | 5                               |         | Running Average Thermal Limit Status (RO)<br>When set, frequency is reduced due to running average thermal limit.                             |
|                  |      | 6                               |         | VR Therm Alert Status (RO)<br>When set, frequency is reduced due to a thermal alert from a processor Voltage Regulator.                       |
|                  |      | 7                               |         | VR Thermal Design Current Status (RO)<br>When set, frequency is reduced due to VR TDC limit.  |
|                  |      | 8                               |         | Other Status (RO)<br>When set, frequency is reduced due to electrical or other constraints.   |
|                  |      | 9                               |         | Reserved  |
|                  |      | 10                              |         | Package/Platform-Level Power Limiting PL1 Status (RO)<br>When set, frequency is reduced due to package/platform-level power limiting PL1.     |
|                  |      | 11                              |         | Package/Platform-Level PL2 Power Limiting Status (RO)<br>When set, frequency is reduced due to package/platform-level power limiting PL2/PL3. |
|                  |      | 12                              |         | Inefficient Operation Status (RO)<br>When set, processor graphics frequency is operating below target frequency.                              |
|                  |      | 15:13                           |         | Reserved  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 20:18                      |       | Reserved.  |
|                  |     | 21                         |       | Running Average Thermal Limit Log<br>When set, indicates that the RATL Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |     | 22                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 23                         |       | VR Thermal Design Current Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.  |
|                  |     | 24                         |       | Other Log<br>When set, indicates that the OTHER Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |     | 25                         |       | Reserved   |
|                  |     | 26                         |       | Package/Platform-Level PL1 Power Limiting Log<br>When set, indicates that the Package/Platform Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description  |
|------------------|------|-----------------------------|---------|--|
| Hex              | Dec  |                             |         |  |
|                  |      | 27                          |         | Package/Platform-Level PL2 Power Limiting Log<br>When set, indicates that the Package/Platform Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |      | 28                          |         | Inefficient Operation Log<br>When set, indicates that the Inefficient Operation Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 63:29                       |         | Reserved   |
| 6B1H             | 1713 | MSR_RING_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in the Ring Interconnect (R/W)<br>(Frequency refers to ring interconnect in the uncore.)   |
|                  |      | 0                           |         | PROCHOT Status (RO)<br>When set, frequency is reduced due to assertion of external PROCHOT.  |
|                  |      | 1                           |         | Thermal Status (RO)<br>When set, frequency is reduced due to a thermal event.  |
|                  |      | 4:2                         |         | Reserved   |
|                  |      | 5                           |         | Running Average Thermal Limit Status (RO)<br>When set, frequency is reduced due to running average thermal limit.  |
|                  |      | 6                           |         | VR Therm Alert Status (RO)<br>When set, frequency is reduced due to a thermal alert from a processor Voltage Regulator.  |
|                  |      | 7                           |         | VR Thermal Design Current Status (RO)<br>When set, frequency is reduced due to VR TDC limit.   |
|                  |      | 8                           |         | Other Status (RO)<br>When set, frequency is reduced due to electrical or other constraints.  |
|                  |      | 9                           |         | Reserved   |
|                  |      | 10                          |         | Package/Platform-Level Power Limiting PL1 Status (RO)<br>When set, frequency is reduced due to package/Platform-level power limiting PL1.  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 11                         |       | Package/Platform-Level PL2 Power Limiting Status (RO)<br>When set, frequency is reduced due to package/Platform-level power limiting PL2/PL3.   |
|                  |     | 15:12                      |       | Reserved  |
|                  |     | 16                         |       | PROCHOT Log<br>When set, indicates that the PROCHOT Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                          |
|                  |     | 17                         |       | Thermal Log<br>When set, indicates that the Thermal Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                          |
|                  |     | 20:18                      |       | Reserved  |
|                  |     | 21                         |       | Running Average Thermal Limit Log<br>When set, indicates that the RATL Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.       |
|                  |     | 22                         |       | VR Therm Alert Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.            |
|                  |     | 23                         |       | VR Thermal Design Current Log<br>When set, indicates that the VR Therm Alert Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0. |
|                  |     | 24                         |       | Other Log<br>When set, indicates that the OTHER Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.                              |
|                  |     | 25                         |       | Reserved  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
|                  |      | 26                         |        | Package/Platform-Level PL1 Power Limiting Log<br>When set, indicates that the Package/Platform Level PL1 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 27                         |        | Package/Platform-Level PL2 Power Limiting Log<br>When set, indicates that the Package/Platform Level PL2 Power Limiting Status bit has asserted since the log bit was last cleared.<br>This log bit will remain set until cleared by software writing 0.   |
|                  |      | 63:28                      |        | Reserved   |
| 6D0H             | 1744 | MSR_LASTBRANCH_16_TO_IP    | Thread | Last Branch Record 16 To IP (R/W)<br>One of 32 triplets of last branch record registers on the last branch record stack. This part of the stack contains pointers to the destination instruction. See also:<br><ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.12.</li> </ul> |
| 6D1H             | 1745 | MSR_LASTBRANCH_17_TO_IP    | Thread | Last Branch Record 17 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D2H             | 1746 | MSR_LASTBRANCH_18_TO_IP    | Thread | Last Branch Record 18 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D3H             | 1747 | MSR_LASTBRANCH_19_TO_IP    | Thread | Last Branch Record 19 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D4H             | 1748 | MSR_LASTBRANCH_20_TO_IP    | Thread | Last Branch Record 20 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D5H             | 1749 | MSR_LASTBRANCH_21_TO_IP    | Thread | Last Branch Record 21 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D6H             | 1750 | MSR_LASTBRANCH_22_TO_IP    | Thread | Last Branch Record 22 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D7H             | 1751 | MSR_LASTBRANCH_23_TO_IP    | Thread | Last Branch Record 23 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D8H             | 1752 | MSR_LASTBRANCH_24_TO_IP    | Thread | Last Branch Record 24 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6D9H             | 1753 | MSR_LASTBRANCH_25_TO_IP    | Thread | Last Branch Record 25 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6DAH             | 1754 | MSR_LASTBRANCH_26_TO_IP    | Thread | Last Branch Record 26 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 6DBH             | 1755 | MSR_LASTBRANCH_27_TO_IP    | Thread  | Last Branch Record 27 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6DCH             | 1756 | MSR_LASTBRANCH_28_TO_IP    | Thread  | Last Branch Record 28 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6DDH             | 1757 | MSR_LASTBRANCH_29_TO_IP    | Thread  | Last Branch Record 29 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6DEH             | 1758 | MSR_LASTBRANCH_30_TO_IP    | Thread  | Last Branch Record 30 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 6DFH             | 1759 | MSR_LASTBRANCH_31_TO_IP    | Thread  | Last Branch Record 31 To IP (R/W)<br>See description of MSR_LASTBRANCH_0_TO_IP.  |
| 770H             | 1904 | IA32_PM_ENABLE             | Package | See Section 14.4.2, "Enabling HWP".  |
| 771H             | 1905 | IA32_HWP_CAPABILITIES      | Thread  | See Section 14.4.3, "HWP Performance Range and Dynamic Capabilities".  |
| 772H             | 1906 | IA32_HWP_REQUEST_PKG       | Package | See Section 14.4.4, "Managing HWP".  |
| 773H             | 1907 | IA32_HWP_INTERRUPT         | Thread  | See Section 14.4.6, "HWP Notifications".   |
| 774H             | 1908 | IA32_HWP_REQUEST           | Thread  | See Section 14.4.4, "Managing HWP".  |
|                  |      | 7:0                        |         | Minimum Performance (R/W)  |
|                  |      | 15:8                       |         | Maximum Performance (R/W)  |
|                  |      | 23:16                      |         | Desired Performance (R/W)  |
|                  |      | 31:24                      |         | Energy/Performance Preference (R/W)  |
|                  |      | 41:32                      |         | Activity Window (R/W)  |
|                  |      | 42                         |         | Package Control (R/W)  |
| 63:43            |      | Reserved                   |         |  |
| 777H             | 1911 | IA32_HWP_STATUS            | Thread  | See Section 14.4.5, "HWP Feedback".  |
| D90H             | 3472 | IA32_BNDCFGS               | Thread  | See Table 2-2.   |
| DA0H             | 3488 | IA32_XSS                   | Thread  | See Table 2-2.   |
| DB0H             | 3504 | IA32_PKG_HDC_CTL           | Package | See Section 14.5.2, "Package level Enabling HDC".  |
| DB1H             | 3505 | IA32_PM_CTL1               | Thread  | See Section 14.5.3, "Logical-Processor Level HDC Control".   |
| DB2H             | 3506 | IA32_THREAD_STALL          | Thread  | See Section 14.5.4.1, "IA32_THREAD_STALL".   |
| DC0H             | 3520 | MSR_LBR_INFO_0             | Thread  | Last Branch Record 0 Additional Information (R/W)<br>One of 32 triplet of last branch record registers on the last branch record stack. This part of the stack contains flag, TSX-related and elapsed cycle information. See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.9.1, "LBR Stack."</li> </ul> |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| DC1H             | 3521 | MSR_LBR_INFO_1             | Thread | Last Branch Record 1 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC2H             | 3522 | MSR_LBR_INFO_2             | Thread | Last Branch Record 2 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC3H             | 3523 | MSR_LBR_INFO_3             | Thread | Last Branch Record 3 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC4H             | 3524 | MSR_LBR_INFO_4             | Thread | Last Branch Record 4 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC5H             | 3525 | MSR_LBR_INFO_5             | Thread | Last Branch Record 5 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC6H             | 3526 | MSR_LBR_INFO_6             | Thread | Last Branch Record 6 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC7H             | 3527 | MSR_LBR_INFO_7             | Thread | Last Branch Record 7 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC8H             | 3528 | MSR_LBR_INFO_8             | Thread | Last Branch Record 8 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DC9H             | 3529 | MSR_LBR_INFO_9             | Thread | Last Branch Record 9 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0.  |
| DCAH             | 3530 | MSR_LBR_INFO_10            | Thread | Last Branch Record 10 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DCBH             | 3531 | MSR_LBR_INFO_11            | Thread | Last Branch Record 11 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DCCH             | 3532 | MSR_LBR_INFO_12            | Thread | Last Branch Record 12 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DCDH             | 3533 | MSR_LBR_INFO_13            | Thread | Last Branch Record 13 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DCEH             | 3534 | MSR_LBR_INFO_14            | Thread | Last Branch Record 14 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DCFH             | 3535 | MSR_LBR_INFO_15            | Thread | Last Branch Record 15 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD0H             | 3536 | MSR_LBR_INFO_16            | Thread | Last Branch Record 16 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD1H             | 3537 | MSR_LBR_INFO_17            | Thread | Last Branch Record 17 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD2H             | 3538 | MSR_LBR_INFO_18            | Thread | Last Branch Record 18 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |

**Table 2-38. Additional MSRs Supported by 6th Generation Intel® Core™ Processors and the Intel® Xeon® Processor Scalable Family Based on Skylake Microarchitecture, 7th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture, 8th Generation Intel® Core™ Processors Based on Coffee Lake Microarchitecture, and Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| DD3H             | 3539 | MSR_LBR_INFO_19            | Thread | Last Branch Record 19 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD4H             | 3520 | MSR_LBR_INFO_20            | Thread | Last Branch Record 20 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD5H             | 3521 | MSR_LBR_INFO_21            | Thread | Last Branch Record 21 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD6H             | 3522 | MSR_LBR_INFO_22            | Thread | Last Branch Record 22 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD7H             | 3523 | MSR_LBR_INFO_23            | Thread | Last Branch Record 23 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD8H             | 3524 | MSR_LBR_INFO_24            | Thread | Last Branch Record 24 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DD9H             | 3525 | MSR_LBR_INFO_25            | Thread | Last Branch Record 25 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDAH             | 3526 | MSR_LBR_INFO_26            | Thread | Last Branch Record 26 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDBH             | 3527 | MSR_LBR_INFO_27            | Thread | Last Branch Record 27 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDCH             | 3528 | MSR_LBR_INFO_28            | Thread | Last Branch Record 28 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDDH             | 3529 | MSR_LBR_INFO_29            | Thread | Last Branch Record 29 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDEH             | 3530 | MSR_LBR_INFO_30            | Thread | Last Branch Record 30 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |
| DDFH             | 3531 | MSR_LBR_INFO_31            | Thread | Last Branch Record 31 Additional Information (R/W)<br>See description of MSR_LBR_INFO_0. |

Table 2-39 lists the MSRs of uncore PMU for Intel processors with CUID DisplayFamily\_DisplayModel signatures of 06\_4EH, 06\_5EH, 06\_8EH, 06\_9EH, and 06\_66H.

**Table 2-39. Uncore PMU MSRs Supported by 6th Generation, 7th Generation, and 8th Generation Intel® Core™ Processors, and Future Intel® Core™ Processors**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description                    |
|------------------|-----|----------------------------|---------|------------------------------------|
| Hex              | Dec |                            |         |                                    |
| 394H             | 916 | MSR_UNC_PERF_FIXED_CTRL    | Package | Uncore Fixed Counter Control (R/W) |
|                  |     | 19:0                       |         | Reserved                           |
|                  |     | 20                         |         | Enable overflow propagation.       |



**Table 2-39. Uncore PMU MSRs Supported by 6th Generation, 7th Generation, and 8th Generation Intel® Core™ Processors, and Future Intel® Core™ Processors**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 21                         |         | Reserved   |
|                  |      | 22                         |         | Enable counting.   |
|                  |      | 63:23                      |         | Reserved   |
| 395H             | 917  | MSR_UNC_PERF_FIXED_CTR     | Package | Uncore Fixed Counter   |
|                  |      | 43:0                       |         | Current count.   |
|                  |      | 63:44                      |         | Reserved   |
| 396H             | 918  | MSR_UNC_CBO_CONFIG         | Package | Uncore C-Box Configuration Information (R/O)   |
|                  |      | 3:0                        |         | Specifies the number of C-Box units with programmable counters (including processor cores and processor graphics). |
|                  |      | 63:4                       |         | Reserved   |
| 3B0H             | 946  | MSR_UNC_ARB_PERFCTR0       | Package | Uncore Arb Unit, Performance Counter 0   |
| 3B1H             | 947  | MSR_UNC_ARB_PERFCTR1       | Package | Uncore Arb Unit, Performance Counter 1   |
| 3B2H             | 944  | MSR_UNC_ARB_PERFEVTSELO    | Package | Uncore Arb Unit, Counter 0 Event Select MSR  |
| 3B3H             | 945  | MSR_UNC_ARB_PERFEVTSEL1    | Package | Uncore Arb Unit, Counter 1 Event Select MSR  |
| 700H             | 1792 | MSR_UNC_CBO_0_PERFEVTSELO  | Package | Uncore C-Box 0, Counter 0 Event Select MSR   |
| 701H             | 1793 | MSR_UNC_CBO_0_PERFEVTSEL1  | Package | Uncore C-Box 0, Counter 1 Event Select MSR   |
| 706H             | 1798 | MSR_UNC_CBO_0_PERFCTR0     | Package | Uncore C-Box 0, Performance Counter 0  |
| 707H             | 1799 | MSR_UNC_CBO_0_PERFCTR1     | Package | Uncore C-Box 0, Performance Counter 1  |
| 710H             | 1808 | MSR_UNC_CBO_1_PERFEVTSELO  | Package | Uncore C-Box 1, Counter 0 Event Select MSR   |
| 711H             | 1809 | MSR_UNC_CBO_1_PERFEVTSEL1  | Package | Uncore C-Box 1, Counter 1 Event Select MSR   |
| 716H             | 1814 | MSR_UNC_CBO_1_PERFCTR0     | Package | Uncore C-Box 1, Performance Counter 0  |
| 717H             | 1815 | MSR_UNC_CBO_1_PERFCTR1     | Package | Uncore C-Box 1, Performance Counter 1  |
| 720H             | 1824 | MSR_UNC_CBO_2_PERFEVTSELO  | Package | Uncore C-Box 2, Counter 0 Event Select MSR   |
| 721H             | 1825 | MSR_UNC_CBO_2_PERFEVTSEL1  | Package | Uncore C-Box 2, Counter 1 Event Select MSR   |
| 726H             | 1830 | MSR_UNC_CBO_2_PERFCTR0     | Package | Uncore C-Box 2, Performance Counter 0  |
| 727H             | 1831 | MSR_UNC_CBO_2_PERFCTR1     | Package | Uncore C-Box 2, Performance Counter 1  |
| 730H             | 1840 | MSR_UNC_CBO_3_PERFEVTSELO  | Package | Uncore C-Box 3, Counter 0 Event Select MSR   |
| 731H             | 1841 | MSR_UNC_CBO_3_PERFEVTSEL1  | Package | Uncore C-Box 3, Counter 1 Event Select MSR   |
| 736H             | 1846 | MSR_UNC_CBO_3_PERFCTR0     | Package | Uncore C-Box 3, Performance Counter 0  |
| 737H             | 1847 | MSR_UNC_CBO_3_PERFCTR1     | Package | Uncore C-Box 3, Performance Counter 1  |
| E01H             | 3585 | MSR_UNC_PERF_GLOBAL_CTRL   | Package | Uncore PMU Global Control  |
|                  |      | 0                          |         | Slice 0 select.  |
|                  |      | 1                          |         | Slice 1 select.  |
|                  |      | 2                          |         | Slice 2 select.  |
|                  |      | 3                          |         | Slice 3 select.  |
|                  |      | 4                          |         | Slice 4select.   |

**Table 2-39. Uncore PMU MSRs Supported by 6th Generation, 7th Generation, and 8th Generation Intel® Core™ Processors, and Future Intel® Core™ Processors**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description                           |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 18:5                       |         | Reserved                                  |
|                  |      | 29                         |         | Enable all uncore counters.               |
|                  |      | 30                         |         | Enable wake on PMI.                       |
|                  |      | 31                         |         | Enable Freezing counter when overflow.    |
|                  |      | 63:32                      |         | Reserved                                  |
| E02H             | 3586 | MSR_UNC_PERF_GLOBAL_STATUS | Package | Uncore PMU Main Status                    |
|                  |      | 0                          |         | Fixed counter overflowed.                 |
|                  |      | 1                          |         | An ARB counter overflowed.                |
|                  |      | 2                          |         | Reserved                                  |
|                  |      | 3                          |         | A CBox counter overflowed (on any slice). |
|                  |      | 63:4                       |         | Reserved                                  |

### 2.16.1 MSRs Specific to 7th Generation and 8th Generation Intel® Core™ Processors based on Kaby Lake Microarchitecture and Coffee Lake Microarchitecture

Table 2-41 lists additional MSRs for 7th generation and 8th generation Intel Core processors with a CPUID DisplayFamily\_DisplayModel signatures of 06\_8EH and 06\_9EH. For an MSR listed in Table 2-41 that also appears in the model-specific tables of prior generations, Table 2-41 supersedes prior generation tables.

**Table 2-40. Additional MSRs Supported by 7th Generation and 8th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture and Coffee Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields     | Scope   | Bit Description  |
|------------------|-----|--------------------------------|---------|--|
| Hex              | Dec |                                |         |  |
| 80H              | 128 | MSR_TRACE_HUB_STH ACPIBAR_BASE | Package | NPK Address Used by AET Messages (R/W)   |
|                  |     | 0                              |         | Lock Bit<br>If set, this MSR cannot be re-written anymore. Lock bit has to be set in order for the AET packets to be directed to NPK MMIO. |
|                  |     | 17:1                           |         | Reserved   |
|                  |     | 63:18                          |         | ACPIBAR_BASE_ADDRESS<br>AET target address in NPK MMIO space.  |
| 1F4H             | 500 | MSR_PRMRR_PHYS_BASE            | Core    | Processor Reserved Memory Range Register - Physical Base Control Register (R/W)  |
|                  |     | 2:0                            |         | MemType<br>PRMRR BASE MemType.   |
|                  |     | 11:3                           |         | Reserved   |
|                  |     | 45:12                          |         | Base<br>PRMRR Base Address.  |
|                  |     | 63:46                          |         | Reserved   |

**Table 2-40. Additional MSRs Supported by 7th Generation and 8th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture and Coffee Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 1F5H             | 501 | MSR_PLMRR_PHYS_MASK        | Core    | Processor Reserved Memory Range Register - Physical Mask Control Register (R/W)  |
|                  |     | 9:0                        |         | Reserved   |
|                  |     | 10                         |         | Lock<br>Lock bit for the PLMRR.  |
|                  |     | 11                         |         | VLD<br>Enable bit for the PLMRR.   |
|                  |     | 45:12                      |         | Mask<br>PLMRR MASK bits.   |
|                  |     | 63:46                      |         | Reserved   |
| 1FBH             | 507 | MSR_PLMRR_VALID_CONFIG     | Core    | Valid PLMRR Configurations (R/W)   |
|                  |     | 0                          |         | 1M supported MEE size.   |
|                  |     | 4:1                        |         | Reserved   |
|                  |     | 5                          |         | 32M supported MEE size.  |
|                  |     | 6                          |         | 64M supported MEE size.  |
|                  |     | 7                          |         | 128M supported MEE size.   |
|                  |     | 31:8                       |         | Reserved   |
| 2F4H             | 756 | MSR_UNCORE_PLMRR_PHYS_BASE | Package | (R/W)<br>The PLMRR range is used to protect Xcode memory from unauthorized reads and writes. Any IO access to this range is aborted. This register controls the location of the PLMRR range by indicating its starting address. It functions in tandem with the PLMRR mask register. |
|                  |     | 11:0                       |         | Reserved   |
|                  |     | 38:12                      |         | Range Base<br>This field corresponds to bits 38:12 of the base address memory range which is allocated to PLMRR memory.  |
|                  |     | 63:39                      |         | Reserved   |
| 2F5H             | 757 | MSR_UNCORE_PLMRR_PHYS_MASK | Package | (R/W)<br>This register controls the size of the PLMRR range by indicating which address bits must match the PLMRR base register value.   |
|                  |     | 9:0                        |         | Reserved   |
|                  |     | 10                         |         | Lock<br>Setting this bit locks all writeable settings in this register, including itself.  |
|                  |     | 11                         |         | Range_En<br>Indicates whether the PLMRR range is enabled and valid.  |

**Table 2-40. Additional MSRs Supported by 7th Generation and 8th Generation Intel® Core™ Processors Based on Kaby Lake Microarchitecture and Coffee Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 38:12                      |         | Range_Mask<br>This field indicates which address bits must match PRMRR base in order to qualify as an PRMRR access. |
|                  |      | 63:39                      |         | Reserved  |
| 620H             | 1568 | MSR_RING_RATIO_LIMIT       | Package | Ring Ratio Limit (R/W)<br>This register provides Min/Max Ratio Limits for the LLC and Ring.                         |
|                  |      | 6:0                        |         | MAX_Ratio<br>This field is used to limit the max ratio of the LLC/Ring.   |
|                  |      | 7                          |         | Reserved  |
|                  |      | 14:8                       |         | MIN_Ratio<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.                             |
|                  |      | 63:15                      |         | Reserved  |

### 2.16.2 MSRs Specific to Future Intel® Core™ Processors

Table 2-41 lists additional MSRs for Future Intel Core processors with a CPUID DisplayFamily\_DisplayModel signature of 06\_66H. For an MSR listed in Table 2-41 that also appears in the model-specific tables of prior generations, Table 2-41 supersedes prior generation tables.

**Table 2-41. Additional MSRs Supported by Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2.   |
|                  |     | 0                          |        | Lock (R/WL)  |
|                  |     | 1                          |        | Enable VMX Inside SMX Operation (R/WL)   |
|                  |     | 2                          |        | Enable VMX Outside SMX Operation (R/WL)  |
|                  |     | 14:8                       |        | SENTER Local Functions Enables (R/WL)  |
|                  |     | 15                         |        | SENTER Global Functions Enable (R/WL)  |
|                  |     | 17                         |        | SGX Launch Control Enable (R/WL)<br>This bit must be set to enable runtime reconfiguration of SGX Launch Control via IA32_SGXLEPUBKEYHASHn MSR.<br>Available only if CPUID.(EAX=07H, ECX=0H): ECX[30] = 1. |
|                  |     | 18                         |        | SGX Global Functions Enable (R/WL)   |
|                  |     | 63:21                      |        | Reserved   |

Table 2-41. Additional MSRs Supported by Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|-----|----------------------------|-------|--|
| Hex              | Dec |                            |       |  |
| 350H             | 848 | MSR_BR_DETECT_CTRL         |       | Branch Monitoring Global Control (R/W)   |
|                  |     | 0                          |       | EnMonitoring<br>Global enable for branch monitoring.   |
|                  |     | 1                          |       | EnExcept<br>Enable branch monitoring event signaling on threshold trip.<br>The branch monitoring event handler is signaled via the existing PMI signaling mechanism as programmed from the corresponding local APIC LVT entry.                                 |
|                  |     | 2                          |       | EnLBRFrz<br>Enable LBR freeze on threshold trip. This will cause the LBR frozen bit 58 to be set in IA32_PERF_GLOBAL_STATUS when a triggering condition occurs and this bit is enabled.  |
|                  |     | 3                          |       | DisableInGuest<br>When set to '1', branch monitoring, event triggering and LBR freeze actions are disabled when operating at VMX non-root operation.   |
|                  |     | 7:4                        |       | Reserved   |
|                  |     | 17:8                       |       | WindowSize<br>Window size defined by WindowCntSel. Values 0 - 1023 are supported.<br>Once the Window counter reaches the WindowSize count both the Window Counter and all Branch Monitoring Counters are cleared.  |
|                  |     | 23:18                      |       | Reserved   |
|                  |     | 25:24                      |       | WindowCntSel<br>Window event count select:<br>'00 = Instructions retired.<br>'01 = Branch instructions retired<br>'10 = Return instructions retired.<br>'11 = Indirect branch instructions retired.  |
|                  |     | 26                         |       | CntAndMode<br>When set to '1', the overall branch monitoring event triggering condition is true only if all enabled counters' threshold conditions are true.<br>When '0', the threshold tripping condition is true if any enabled counters' threshold is true. |
| 63:27            |     | Reserved                   |       |  |
| 351H             | 849 | MSR_BR_DETECT_STATUS       |       | Branch Monitoring Global Status (R/W)  |
|                  |     | 0                          |       | Branch Monitoring Event Signaled<br>When set to '1', Branch Monitoring event signaling is blocked until this bit is cleared by software.   |

**Table 2-41. Additional MSRs Supported by Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |     | Register Name / Bit Fields | Scope | Bit Description   |
|------------------|-----|----------------------------|-------|---|
| Hex              | Dec |                            |       |   |
|                  |     | 1                          |       | LBRsValid<br>This status bit is set to '1' if the LBR state is considered valid for sampling by branch monitoring software.   |
|                  |     | 7:2                        |       | Reserved  |
|                  |     | 8                          |       | CntrHit0<br>Branch monitoring counter #0 threshold hit. This status bit is sticky and once set requires clearing by software. Counter operation continues independent of the state of the bit.  |
|                  |     | 9                          |       | CntrHit1<br>Branch monitoring counter #1 threshold hit. This status bit is sticky and once set requires clearing by software. Counter operation continues independent of the state of the bit.  |
|                  |     | 15:10                      |       | Reserved<br>Reserved for additional branch monitoring counters threshold hit status.  |
|                  |     | 25:16                      |       | CountWindow<br>The current value of the window counter. The count value is frozen on a valid branch monitoring triggering condition. This is a 10-bit unsigned value.   |
|                  |     | 31:26                      |       | Reserved<br>Reserved for future extension of CountWindow.   |
|                  |     | 39:32                      |       | Count0<br>The current value of counter 0 updated after each occurrence of the event being counted. The count value is frozen on a valid branch monitoring triggering condition (in which case CntrHit0 will also be set). This is an 8-bit signed value (2's complement).<br>Heuristic events which only increment will saturate and freeze at maximum value 0xFF (256).<br>RET-CALL event counter saturate at maximum value 0x7F (+127) and minimum value 0x80 (-128). |
|                  |     | 47:40                      |       | Count1<br>The current value of counter 1 updated after each occurrence of the event being counted. The count value is frozen on a valid branch monitoring triggering condition (in which case CntrHit1 will also be set). This is an 8-bit signed value (2's complement).<br>Heuristic events which only increment will saturate and freeze at maximum value 0xFF (256).<br>RET-CALL event counter saturate at maximum value 0x7F (+127) and minimum value 0x80 (-128). |
|                  |     | 63:48                      |       | Reserved  |

Table 2-41. Additional MSRs Supported by Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture

| Register Address  |                 | Register Name / Bit Fields     | Scope   | Bit Description   |
|-------------------|-----------------|--------------------------------|---------|---|
| Hex               | Dec             |                                |         |   |
| 354H<br>-<br>355H | 852<br>-<br>853 | MSR_BR_DETECT_COUNTER_CONFIG_i |         | Branch Monitoring Detect Counter Configuration (R/W)  |
|                   |                 | 0                              |         | CntrEn<br>Enable counter.   |
|                   |                 | 7:1                            |         | CntrEvSel<br>Event select (other values #GP)<br>'0000000 = RETs.<br>'0000001 = RET-CALL bias.<br>'0000010 = RET mispredicts.<br>'0000011 = Branch (all) mispredicts.<br>'0000100 = Indirect branch mispredicts.<br>'0000101 = Far branch instructions.  |
|                   |                 | 14:8                           |         | CntrThreshold<br>Threshold (an unsigned value of 0 to 127 supported).<br>The value 0 of counter threshold will result in event signaled after every instruction. #GP if threshold is < 2.   |
|                   |                 | 15                             |         | MispredEventCnt<br>Mispredict events counting behavior:<br>'0 = Mispredict events are counted in a window.<br>'1 = Mispredict events are counted based on a consecutive occurrence. CntrThreshold is treated as # of consecutive mispredicts. This control bit only applies to events specified by CntrEvSel that involve a prediction (0000010, 0000011, 0000100). Setting this bit for other events is ignored. |
|                   |                 | 63:16                          |         | Reserved  |
| 3F8H              | 1016            | MSR_PKG_C3_RESIDENCY           | Package | Package C3 Residency Counter (R/O)  |
|                   |                 | 63:0                           |         | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |
| 620H              | 1568            | MSR_RING_RATIO_LIMIT           | Package | Ring Ratio Limit (R/W)<br>This register provides Min/Max Ratio Limits for the LLC and Ring.   |
|                   |                 | 6:0                            |         | MAX_Ratio<br>This field is used to limit the max ratio of the LLC/Ring.   |
|                   |                 | 7                              |         | Reserved  |
|                   |                 | 14:8                           |         | MIN_Ratio<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.   |
|                   |                 | 63:15                          |         | Reserved  |

**Table 2-41. Additional MSRs Supported by Future Intel® Core™ Processors Based on Cannon Lake Microarchitecture**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|------|----------------------------|-------|--|
| Hex              | Dec  |                            |       |  |
| 660H             | 1632 | MSR_CORE_C1_RESIDENCY      | Core  | Core C1 Residency Counter (R/O)  |
|                  |      | 63:0                       |       | Value since last reset for the Core C1 residency. Counter rate is the Max Non-Turbo frequency (same as TSC). This counter counts in case both of the core's threads are in an idle state and at least one of the core's thread residency is in a C1 state or in one of its sub states. The counter is updated only after a core C state exit. Note: Always reads 0 if core C1 is unsupported. A value of zero indicates that this processor does not support core C1 or never entered core C1 level state. |
| 662H             | 1634 | MSR_CORE_C3_RESIDENCY      | Core  | Core C3 Residency Counter (R/O)  |
|                  |      | 63:0                       |       | Will always return 0.  |

### 2.16.3 MSRs Specific to Intel® Xeon® Processor Scalable Family

Intel® Xeon® Processor Scalable Family (CPUID DisplayFamily\_DisplayModel = 06\_55H) support the MSRs listed in Table 2-42.

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |       | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-------|----------------------------|---------|--|
| Hex              | Dec   |                            |         |  |
| 3AH              | 58    | IA32_FEATURE_CONTROL       | Thread  | Control Features in Intel 64 Processor (R/W)<br>See Table 2-2. |
|                  |       | 0                          |         | Lock (R/WL)  |
|                  |       | 1                          |         | Enable VMX Inside SMX Operation (R/WL)                         |
|                  |       | 2                          |         | Enable VMX Outside SMX Operation (R/WL)                        |
|                  |       | 14:8                       |         | SENTER Local Functions Enables (R/WL)                          |
|                  |       | 15                         |         | SENTER Global Functions Enable (R/WL)                          |
|                  |       | 18                         |         | SGX Global Functions Enable (R/WL)                             |
|                  |       | 20                         |         | LMCE_ON (R/WL)   |
|                  | 63:21 | Reserved                   |         |  |
| 4EH              | 78    | MSR_PPIN_CTL               | Package | Protected Processor Inventory Number Enable Control (R/W)      |
|                  |       | 0                          |         | LockOut (R/WO)<br>See Table 2-25.                              |
|                  |       | 1                          |         | Enable_PPIN (R/W)<br>See Table 2-25.                           |
|                  |       | 63:2                       |         | Reserved   |
| 4FH              | 79    | MSR_PPIN                   | Package | Protected Processor Inventory Number (R/O)                     |



Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 63:0                       |         | Protected Processor Inventory Number (R/O)<br>See Table 2-25.  |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .   |
|                  |     | 7:0                        |         | Reserved   |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>See Table 2-25.   |
|                  |     | 22:16                      |         | Reserved.  |
|                  |     | 23                         | Package | PPIN_CAP (R/O)<br>See Table 2-25.  |
|                  |     | 27:24                      |         | Reserved   |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>See Table 2-25.   |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>See Table 2-25.   |
|                  |     | 30                         | Package | Programmable TJ OFFSET (R/O)<br>See Table 2-25.  |
|                  |     | 39:31                      |         | Reserved   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>See Table 2-25.  |
|                  |     | 63:48                      |         | Reserved   |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Core    | C-State Configuration Control (R/W)<br>Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.<br>See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest processor-specific C-state code name (consuming the least power) for the package. The default is set as factory-configured package C-state limit.<br>The following C-state code name encodings are supported:<br>000b: C0/C1 (no package C-state support)<br>001b: C2<br>010b: C6 (non-retention)<br>011b: C6 (retention)<br>111b: No Package C state limits. All C states supported by the processor are available. |
|                  |     | 9:3                        |         | Reserved   |

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
|                  |     | 10                         |        | I/O MWAIT Redirection Enable (R/W)  |
|                  |     | 14:11                      |        | Reserved  |
|                  |     | 15                         |        | CFG Lock (R/WO)   |
|                  |     | 16                         |        | Automatic C-State Conversion Enable (R/W)<br>If 1, the processor will convert HALT or MWAIT(C1) to MWAIT(C6).   |
|                  |     | 24:17                      |        | Reserved  |
|                  |     | 25                         |        | C3 State Auto Demotion Enable (R/W)   |
|                  |     | 26                         |        | C1 State Auto Demotion Enable (R/W)   |
|                  |     | 27                         |        | Enable C3 Undemotion (R/W)  |
|                  |     | 28                         |        | Enable C1 Undemotion (R/W)  |
|                  |     | 29                         |        | Package C State Demotion Enable (R/W)   |
|                  |     | 30                         |        | Package C State UnDemotion Enable (R/W)   |
|                  |     | 63:31                      |        | Reserved  |
| 179H             | 377 | IA32_MCG_CAP               | Thread | Global Machine Check Capability (R/O)   |
|                  |     | 7:0                        |        | Count   |
|                  |     | 8                          |        | MCG_CTL_P   |
|                  |     | 9                          |        | MCG_EXT_P   |
|                  |     | 10                         |        | MCP_CMCI_P  |
|                  |     | 11                         |        | MCG_TES_P   |
|                  |     | 15:12                      |        | Reserved  |
|                  |     | 23:16                      |        | MCG_EXT_CNT   |
|                  |     | 24                         |        | MCG_SER_P   |
|                  |     | 25                         |        | MCG_EM_P  |
|                  |     | 26                         |        | MCG_ELOG_P  |
|                  |     | 63:27                      |        | Reserved  |
| 17DH             | 390 | MSR_SMM_MCA_CAP            | THREAD | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability Enhancement. Accessible only while in SMM.   |
|                  |     | 57:0                       |        | Reserved  |
|                  |     | 58                         |        | SMM_Code_Access_Chk (SMM-RO)<br>If set to 1 indicates that the SMM code access restriction is supported and a host-space interface is available to SMM handler. |
|                  |     | 59                         |        | Long_Flow_Indication (SMM-RO)<br>If set to 1 indicates that the SMM long flow indicator is supported and a host-space interface is available to SMM handler.    |
|                  |     | 63:60                      |        | Reserved  |

Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |     | Register Name / Bit Fields                           | Scope | Bit Description   |
|------------------|-----|--|-------|---|
| Hex              | Dec |  |       |   |
| 19CH             | 412 | IA32_THERM_STATUS                                    | Core  | Thermal Monitor Status (R/W)<br>See Table 2-2.            |
|                  |     | 0  |       | Thermal Status (RO)<br>See Table 2-2.                     |
|                  |     | 1  |       | Thermal Status Log (R/WCO)<br>See Table 2-2.              |
|                  |     | 2  |       | PROTCHOT # or FORCEPR# Status (RO)<br>See Table 2-2.      |
|                  |     | 3  |       | PROTCHOT # or FORCEPR# Log (R/WCO)<br>See Table 2-2.      |
|                  |     | 4  |       | Critical Temperature Status (RO)<br>See Table 2-2.        |
|                  |     | 5  |       | Critical Temperature Status Log (R/WCO)<br>See Table 2-2. |
|                  |     | 6  |       | Thermal Threshold #1 Status (RO)<br>See Table 2-2.        |
|                  |     | 7  |       | Thermal Threshold #1 Log (R/WCO)<br>See Table 2-2.        |
|                  |     | 8  |       | Thermal Threshold #2 Status (RO)<br>See Table 2-2.        |
|                  |     | 9  |       | Thermal Threshold #2 Log (R/WCO)<br>See Table 2-2.        |
|                  |     | 10   |       | Power Limitation Status (RO)<br>See Table 2-2.            |
|                  |     | 11   |       | Power Limitation Log (R/WCO)<br>See Table 2-2.            |
|                  |     | 12   |       | Current Limit Status (RO)<br>See Table 2-2.               |
|                  |     | 13   |       | Current Limit Log (R/WCO)<br>See Table 2-2.               |
|                  |     | 14   |       | Cross Domain Limit Status (RO)<br>See Table 2-2.          |
|                  |     | 15   |       | Cross Domain Limit Log (R/WCO)<br>See Table 2-2.          |
|                  |     | 22:16  |       | Digital Readout (RO)<br>See Table 2-2.                    |
| 26:23            |     | Reserved   |       |   |
| 30:27            |     | Resolution in Degrees Celsius (RO)<br>See Table 2-2. |       |   |

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 31                         |         | Reading Valid (RO)<br>See Table 2-2.  |
|                  |     | 63:32                      |         | Reserved  |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Package | Temperature Target  |
|                  |     | 15:0                       |         | Reserved  |
|                  |     | 23:16                      |         | Temperature Target (RO)<br>See Table 2-25.  |
|                  |     | 27:24                      |         | TCC Activation Offset (R/W)<br>See Table 2-25.  |
|                  |     | 63:28                      |         | Reserved  |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | This register defines the ratio limits. RATIO[0:7] must be populated in ascending order. RATIO[i+1] must be less than or equal to RATIO[i]. Entries with RATIO[i] will be ignored. If any of the rules above are broken, the configuration is silently rejected. If the programmed ratio is: <ul style="list-style-type: none"> <li>▪ Above the fused ratio for that core count, it will be clipped to the fuse limits (assuming !OC).</li> <li>▪ Below the min supported ratio, it will be clipped.</li> </ul> |
|                  |     | 7:0                        |         | RATIO_0<br>Defines ratio limits.  |
|                  |     | 15:8                       |         | RATIO_1<br>Defines ratio limits.  |
|                  |     | 23:16                      |         | RATIO_2<br>Defines ratio limits.  |
|                  |     | 31:24                      |         | RATIO_3<br>Defines ratio limits.  |
|                  |     | 39:32                      |         | RATIO_4<br>Defines ratio limits.  |
|                  |     | 47:40                      |         | RATIO_5<br>Defines ratio limits.  |
|                  |     | 55:48                      |         | RATIO_6<br>Defines ratio limits.  |
|                  |     | 63:56                      |         | RATIO_7<br>Defines ratio limits.  |
|                  |     | 1AEH                       | 430     | MSR_TURBO_RATIO_LIMIT_CORES   |

Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 7:0                        |         | NUMCORE_0<br>Defines the active core ranges for each frequency point. |
|                  |     | 15:8                       |         | NUMCORE_1<br>Defines the active core ranges for each frequency point. |
|                  |     | 23:16                      |         | NUMCORE_2<br>Defines the active core ranges for each frequency point. |
|                  |     | 31:24                      |         | NUMCORE_3<br>Defines the active core ranges for each frequency point. |
|                  |     | 39:32                      |         | NUMCORE_4<br>Defines the active core ranges for each frequency point. |
|                  |     | 47:40                      |         | NUMCORE_5<br>Defines the active core ranges for each frequency point. |
|                  |     | 55:48                      |         | NUMCORE_6<br>Defines the active core ranges for each frequency point. |
|                  |     | 63:56                      |         | NUMCORE_7<br>Defines the active core ranges for each frequency point. |
| 280H             | 640 | IA32_MC0_CTL2              | Core    | See Table 2-2.  |
| 281H             | 641 | IA32_MC1_CTL2              | Core    | See Table 2-2.  |
| 282H             | 642 | IA32_MC2_CTL2              | Core    | See Table 2-2.  |
| 283H             | 643 | IA32_MC3_CTL2              | Core    | See Table 2-2.  |
| 284H             | 644 | IA32_MC4_CTL2              | Package | See Table 2-2.  |
| 285H             | 645 | IA32_MC5_CTL2              | Package | See Table 2-2.  |
| 286H             | 646 | IA32_MC6_CTL2              | Package | See Table 2-2.  |
| 287H             | 647 | IA32_MC7_CTL2              | Package | See Table 2-2.  |
| 288H             | 648 | IA32_MC8_CTL2              | Package | See Table 2-2.  |
| 289H             | 649 | IA32_MC9_CTL2              | Package | See Table 2-2.  |
| 28AH             | 650 | IA32_MC10_CTL2             | Package | See Table 2-2.  |
| 28BH             | 651 | IA32_MC11_CTL2             | Package | See Table 2-2.  |
| 28CH             | 652 | IA32_MC12_CTL2             | Package | See Table 2-2.  |
| 28DH             | 653 | IA32_MC13_CTL2             | Package | See Table 2-2.  |
| 28EH             | 654 | IA32_MC14_CTL2             | Package | See Table 2-2.  |
| 28FH             | 655 | IA32_MC15_CTL2             | Package | See Table 2-2.  |

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 290H             | 656  | IA32_MC16_CTL2             | Package | See Table 2-2.   |
| 291H             | 657  | IA32_MC17_CTL2             | Package | See Table 2-2.   |
| 292H             | 658  | IA32_MC18_CTL2             | Package | See Table 2-2.   |
| 293H             | 659  | IA32_MC19_CTL2             | Package | See Table 2-2.   |
| 400H             | 1024 | IA32_MCO_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MCO reports MC errors from the IFU module.             |
| 401H             | 1025 | IA32_MCO_STATUS            | Core    |  |
| 402H             | 1026 | IA32_MCO_ADDR              | Core    |  |
| 403H             | 1027 | IA32_MCO_MISC              | Core    |  |
| 404H             | 1028 | IA32_MC1_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC1 reports MC errors from the DCU module.             |
| 405H             | 1029 | IA32_MC1_STATUS            | Core    |  |
| 406H             | 1030 | IA32_MC1_ADDR              | Core    |  |
| 407H             | 1031 | IA32_MC1_MISC              | Core    |  |
| 408H             | 1032 | IA32_MC2_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC2 reports MC errors from the DTLB module.            |
| 409H             | 1033 | IA32_MC2_STATUS            | Core    |  |
| 40AH             | 1034 | IA32_MC2_ADDR              | Core    |  |
| 40BH             | 1035 | IA32_MC2_MISC              | Core    |  |
| 40CH             | 1036 | IA32_MC3_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC3 reports MC errors from the MLC module.             |
| 40DH             | 1037 | IA32_MC3_STATUS            | Core    |  |
| 40EH             | 1038 | IA32_MC3_ADDR              | Core    |  |
| 40FH             | 1039 | IA32_MC3_MISC              | Core    |  |
| 410H             | 1040 | IA32_MC4_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC4 reports MC errors from the PCU module.             |
| 411H             | 1041 | IA32_MC4_STATUS            | Package |  |
| 412H             | 1042 | IA32_MC4_ADDR              | Package |  |
| 413H             | 1043 | IA32_MC4_MISC              | Package |  |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC5 reports MC errors from a link interconnect module. |
| 415H             | 1045 | IA32_MC5_STATUS            | Package |  |
| 416H             | 1046 | IA32_MC5_ADDR              | Package |  |
| 417H             | 1047 | IA32_MC5_MISC              | Package |  |
| 418H             | 1048 | IA32_MC6_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC6 reports MC errors from the integrated I/O module.  |
| 419H             | 1049 | IA32_MC6_STATUS            | Package |  |
| 41AH             | 1050 | IA32_MC6_ADDR              | Package |  |
| 41BH             | 1051 | IA32_MC6_MISC              | Package |  |
| 41CH             | 1052 | IA32_MC7_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC7 reports MC errors from the M2M 0.                  |
| 41DH             | 1053 | IA32_MC7_STATUS            | Package |  |
| 41EH             | 1054 | IA32_MC7_ADDR              | Package |  |
| 41FH             | 1055 | IA32_MC7_MISC              | Package |  |

Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
| 420H             | 1056 | IA32_MC8_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC8 reports MC errors from the M2M 1.  |
| 421H             | 1057 | IA32_MC8_STATUS            | Package |  |
| 422H             | 1058 | IA32_MC8_ADDR              | Package |  |
| 423H             | 1059 | IA32_MC8_MISC              | Package |  |
| 424H             | 1060 | IA32_MC9_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 - MC11 report MC errors from the CHA                                    |
| 425H             | 1061 | IA32_MC9_STATUS            | Package |  |
| 426H             | 1062 | IA32_MC9_ADDR              | Package |  |
| 427H             | 1063 | IA32_MC9_MISC              | Package |  |
| 428H             | 1064 | IA32_MC10_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 - MC11 report MC errors from the CHA.                                   |
| 429H             | 1065 | IA32_MC10_STATUS           | Package |  |
| 42AH             | 1066 | IA32_MC10_ADDR             | Package |  |
| 42BH             | 1067 | IA32_MC10_MISC             | Package |  |
| 42CH             | 1068 | IA32_MC11_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC9 - MC11 report MC errors from the CHA.                                   |
| 42DH             | 1069 | IA32_MC11_STATUS           | Package |  |
| 42EH             | 1070 | IA32_MC11_ADDR             | Package |  |
| 42FH             | 1071 | IA32_MC11_MISC             | Package |  |
| 430H             | 1072 | IA32_MC12_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC12 report MC errors from each channel of a link interconnect module.      |
| 431H             | 1073 | IA32_MC12_STATUS           | Package |  |
| 432H             | 1074 | IA32_MC12_ADDR             | Package |  |
| 433H             | 1075 | IA32_MC12_MISC             | Package |  |
| 434H             | 1076 | IA32_MC13_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers. |
| 435H             | 1077 | IA32_MC13_STATUS           | Package |  |
| 436H             | 1078 | IA32_MC13_ADDR             | Package |  |
| 437H             | 1079 | IA32_MC13_MISC             | Package |  |
| 438H             | 1080 | IA32_MC14_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers. |
| 439H             | 1081 | IA32_MC14_STATUS           | Package |  |
| 43AH             | 1082 | IA32_MC14_ADDR             | Package |  |
| 43BH             | 1083 | IA32_MC14_MISC             | Package |  |
| 43CH             | 1084 | IA32_MC15_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers. |
| 43DH             | 1085 | IA32_MC15_STATUS           | Package |  |
| 43EH             | 1086 | IA32_MC15_ADDR             | Package |  |
| 43FH             | 1087 | IA32_MC15_MISC             | Package |  |
| 440H             | 1088 | IA32_MC16_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers  |
| 441H             | 1089 | IA32_MC16_STATUS           | Package |  |
| 442H             | 1090 | IA32_MC16_ADDR             | Package |  |
| 443H             | 1091 | IA32_MC16_MISC             | Package |  |

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 444H             | 1092 | IA32_MC17_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers.                                      |
| 445H             | 1093 | IA32_MC17_STATUS           | Package |   |
| 446H             | 1094 | IA32_MC17_ADDR             | Package |   |
| 447H             | 1095 | IA32_MC17_MISC             | Package |   |
| 448H             | 1096 | IA32_MC18_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Banks MC13 through MC 18 report MC errors from the integrated memory controllers.                                      |
| 449H             | 1097 | IA32_MC18_STATUS           | Package |   |
| 44AH             | 1098 | IA32_MC18_ADDR             | Package |   |
| 44BH             | 1099 | IA32_MC18_MISC             | Package |   |
| 44CH             | 1100 | IA32_MC19_CTL              | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs" through Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>Bank MC19 reports MC errors from a link interconnect module.   |
| 44DH             | 1101 | IA32_MC19_STATUS           | Package |   |
| 44EH             | 1102 | IA32_MC19_ADDR             | Package |   |
| 44FH             | 1103 | IA32_MC19_MISC             | Package |   |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers Used in RAPL Interfaces (R/O)  |
|                  |      | 3:0                        | Package | Power Units<br>See Section 14.9.1, "RAPL Interfaces."   |
|                  |      | 7:4                        | Package | Reserved  |
|                  |      | 12:8                       | Package | Energy Status Units<br>Energy related information (in Joules) is based on the multiplier, $1/2^{\wedge}ESU$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 0EH (or 61 micro-joules). |
|                  |      | 15:13                      | Package | Reserved  |
|                  |      | 19:16                      | Package | Time Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 63:20                      |         | Reserved  |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>Energy consumed by DRAM devices.  |
|                  |      | 31:0                       |         | Energy in 15.3 micro-joules. Requires BIOS configuration to enable DRAM RAPL mode 0 (Direct VR).  |
|                  |      | 63:32                      |         | Reserved  |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |



Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |           | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----------|----------------------------|---------|---|
| Hex              | Dec       |                            |         |   |
| 620H             | 1568      | MSR_UNCORE_RATIO_LIMIT     | Package | Uncore Ratio Limit (R/W)<br>Out of reset, the min_ratio and max_ratio fields represent the widest possible range of uncore frequencies. Writing to these fields allows software to control the minimum and the maximum frequency that hardware will select. |
|                  |           | 63:15                      |         | Reserved  |
|                  |           | 14:8                       |         | MIN_RATIO<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.   |
|                  |           | 7                          |         | Reserved  |
|                  |           | 6:0                        |         | MAX_RATIO<br>This field is used to limit the max ratio of the LLC/Ring.   |
| 639H             | 1593      | MSR_PPO_ENERGY_STATUS      | Package | Reserved (R/O)<br>Reads return 0.   |
| C8DH             | 3213      | IA32_QM_EVTSEL             | THREAD  | Monitoring Event Select Register (R/W)<br>If CPUID.(EAX=07H, ECX=0):EBX.RDT-M[bit 12] = 1.  |
|                  |           | 7:0                        |         | EventID (RW)<br>Event encoding:<br>0x00: No monitoring.<br>0x01: L3 occupancy monitoring.<br>0x02: Total memory bandwidth monitoring.<br>0x03: Local memory bandwidth monitoring.<br>All other encoding reserved.   |
|                  |           | 31:8                       |         | Reserved  |
|                  |           | 41:32                      |         | RMID (RW)   |
|                  |           | 63:42                      |         | Reserved  |
|                  |           | C8FH                       |         | 3215  |
| 9:0              | RMID      |                            |         |   |
| 31:10            | Reserved  |                            |         |   |
| 51:32            | COS (R/W) |                            |         |   |
| 63: 52           | Reserved  |                            |         |   |
| C90H             | 3216      | IA32_L3_QOS_MASK_0         | Package | L3 Class Of Service Mask - COS 0 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=0.  |
|                  |           | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 0 enforcement.   |
|                  |           | 63:20                      |         | Reserved  |
| C91H             | 3217      | IA32_L3_QOS_MASK_1         | Package | L3 Class Of Service Mask - COS 1 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=1.  |
|                  |           | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 1 enforcement.   |

**Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:20                      |         | Reserved  |
| C92H             | 3218 | IA32_L3_QOS_MASK_2         | Package | L3 Class Of Service Mask - COS 2 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=2.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 2 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C93H             | 3219 | IA32_L3_QOS_MASK_3         | Package | L3 Class Of Service Mask - COS 3 (R/W).<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=3. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 3 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C94H             | 3220 | IA32_L3_QOS_MASK_4         | Package | L3 Class Of Service Mask - COS 4 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=4.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 4 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C95H             | 3221 | IA32_L3_QOS_MASK_5         | Package | L3 Class Of Service Mask - COS 5 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=5.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 5 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C96H             | 3222 | IA32_L3_QOS_MASK_6         | Package | L3 Class Of Service Mask - COS 6 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=6.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 6 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C97H             | 3223 | IA32_L3_QOS_MASK_7         | Package | L3 Class Of Service Mask - COS 7 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=7.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 7 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C98H             | 3224 | IA32_L3_QOS_MASK_8         | Package | L3 Class Of Service Mask - COS 8 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=8.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 8 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved  |
| C99H             | 3225 | IA32_L3_QOS_MASK_9         | Package | L3 Class Of Service Mask - COS 9 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=9.  |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 9 enforcement.                                 |

Table 2-42. MSRs Supported by Intel® Xeon® Processor Scalable Family with DisplayFamily\_DisplayModel 06\_55H

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|------|----------------------------|---------|--|
| Hex              | Dec  |                            |         |  |
|                  |      | 63:20                      |         | Reserved   |
| C9AH             | 3226 | IA32_L3_QOS_MASK_10        | Package | L3 Class Of Service Mask - COS 10 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=10. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 10 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9BH             | 3227 | IA32_L3_QOS_MASK_11        | Package | L3 Class Of Service Mask - COS 11 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=11. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 11 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9CH             | 3228 | IA32_L3_QOS_MASK_12        | Package | L3 Class Of Service Mask - COS 12 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=12. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 12 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9DH             | 3229 | IA32_L3_QOS_MASK_13        | Package | L3 Class Of Service Mask - COS 13 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=13. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 13 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9EH             | 3230 | IA32_L3_QOS_MASK_14        | Package | L3 Class Of Service Mask - COS 14 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=14. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 14 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |
| C9FH             | 3231 | IA32_L3_QOS_MASK_15        | Package | L3 Class Of Service Mask - COS 15 (R/W)<br>If CPUID.(EAX=10H, ECX=1):EDX.COS_MAX[15:0] >=15. |
|                  |      | 0:19                       |         | CBM: Bit vector of available L3 ways for COS 15 enforcement.                                 |
|                  |      | 63:20                      |         | Reserved   |

## 2.17 MSRS IN INTEL® XEON PHI™ PROCESSOR 3200/5200/7200 SERIES AND INTEL® XEON PHI™ PROCESSOR 7215/7285/7295 SERIES

Intel® Xeon Phi™ processor 3200, 5200, 7200 series, with CPUID DisplayFamily\_DisplayModel signature 06\_57H, supports the MSR interfaces listed in Table 2-43. These processors are based on the Knights Landing microarchitecture. Intel® Xeon Phi™ processor 7215, 7285, 7295 series, with CPUID DisplayFamily\_DisplayModel signature 06\_85H, supports the MSR interfaces listed in Table 2-43 and Table 2-44. These processors are based on the Knights Mill microarchitecture. Some MSRs are shared between a pair of processor cores, the scope is marked as module.

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
| 0H               | 0   | IA32_P5_MC_ADDR            | Module  | See Section 2.22, "MSRs in Pentium Processors."  |
| 1H               | 1   | IA32_P5_MC_TYPE            | Module  | See Section 2.22, "MSRs in Pentium Processors."  |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE   | Thread  | See Section 8.10.5, "Monitor/Mwait Address Range Determination." See Table 2-2.  |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER    | Thread  | See Section 17.17, "Time-Stamp Counter," and see Table 2-2.  |
| 17H              | 23  | IA32_PLATFORM_ID           | Package | Platform ID (R)<br>See Table 2-2.  |
| 1BH              | 27  | IA32_APIC_BASE             | Thread  | See Section 10.4.4, "Local APIC Status and Location," and Table 2-2.   |
| 34H              | 52  | MSR_SMI_COUNT              | Thread  | SMI Counter (R/O)  |
|                  |     | 31:0                       |         | SMI Count (R/O)  |
|                  |     | 63:32                      |         | Reserved   |
| 3AH              | 58  | IA32_FEATURE_CONTROL       | Thread  | Control Features in Intel 64Processor (R/W)<br>See Table 2-2.  |
|                  |     | 0                          |         | Lock (R/WL)  |
|                  |     | 1                          |         | Reserved   |
|                  |     | 2                          |         | Enable VMX outside SMX operation (R/WL)  |
| 3BH              | 59  | IA32_TSC_ADJUST            | THREAD  | Per-Logical-Processor TSC ADJUST (R/W)<br>See Table 2-2.   |
| 4EH              | 78  | MSR_PPIN_CTL               | Package | Protected Processor Inventory Number Enable Control (R/W)  |
|                  |     | 0                          |         | LockOut (R/WO)<br>Set 1 to prevent further writes to MSR_PPIN_CTL. Writing 1 to MSR_PPIN_CTL[bit 0] is permitted only if MSR_PPIN_CTL[bit 1] is clear. Default is 0.<br>BIOS should provide an opt-in menu to enable the user to turn on MSR_PPIN_CTL[bit 1] for a privileged inventory initialization agent to access MSR_PPIN. After reading MSR_PPIN, the privileged inventory initialization agent should write '01b' to MSR_PPIN_CTL to disable further access to MSR_PPIN and prevent unauthorized modification to MSR_PPIN_CTL. |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 1                          |         | Enable_PPIN (R/W)<br>If 1, enables MSR_PPIN to be accessible using RDMSR. Once set, an attempt to write 1 to MSR_PPIN_CTL[bit 0] will cause #GP.<br>If 0, an attempt to read MSR_PPIN will cause #GP. Default is 0.   |
|                  |     | 63:2                       |         | Reserved  |
| 4FH              | 79  | MSR_PPIN                   | Package | Protected Processor Inventory Number (R/O)  |
|                  |     | 63:0                       |         | Protected Processor Inventory Number (R/O)<br>A unique value within a given CUID family/model/stepping signature that a privileged inventory initialization agent can access to identify each physical processor, when access to MSR_PPIN is enabled. Access to MSR_PPIN is permitted only if MSR_PPIN_CTL[bits 1:0] = '10b'. |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG        | Core    | BIOS Update Trigger Register (W)<br>See Table 2-2.  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID          | THREAD  | BIOS Update Signature ID (RO)<br>See Table 2-2.   |
| C1H              | 193 | IA32_PMC0                  | THREAD  | Performance Counter Register<br>See Table 2-2.  |
| C2H              | 194 | IA32_PMC1                  | THREAD  | Performance Counter Register<br>See Table 2-2.  |
| CEH              | 206 | MSR_PLATFORM_INFO          | Package | Platform Information<br>Contains power management and other model specific features enumeration. See <a href="http://biosbits.org">http://biosbits.org</a> .  |
|                  |     | 7:0                        |         | Reserved  |
|                  |     | 15:8                       | Package | Maximum Non-Turbo Ratio (R/O)<br>This is the ratio of the frequency that invariant TSC runs at. Frequency = ratio * 100 MHz.  |
|                  |     | 27:16                      |         | Reserved  |
|                  |     | 28                         | Package | Programmable Ratio Limit for Turbo Mode (R/O)<br>When set to 1, indicates that Programmable Ratio Limit for Turbo mode is enabled. When set to 0, indicates Programmable Ratio Limit for Turbo mode is disabled.  |
|                  |     | 29                         | Package | Programmable TDP Limit for Turbo Mode (R/O)<br>When set to 1, indicates that TDP Limit for Turbo mode is programmable. When set to 0, indicates TDP Limit for Turbo mode is not programmable.   |
|                  |     | 39:30                      |         | Reserved  |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 47:40                      | Package | Maximum Efficiency Ratio (R/O)<br>This is the minimum ratio (maximum efficiency) that the processor can operate, in units of 100MHz.  |
|                  |     | 63:48                      |         | Reserved  |
| E2H              | 226 | MSR_PKG_CST_CONFIG_CONTROL | Package | C-State Configuration Control (R/W)   |
|                  |     | 2:0                        |         | Package C-State Limit (R/W)<br>Specifies the lowest C-state for the package. This feature does not limit the processor core C-state. The power-on default value from bit[2:0] of this register reports the deepest package C-state the processor is capable to support when manufactured. It is recommended that BIOS always read the power-on default value reported from this bit field to determine the supported deepest C-state on the processor and leave it as default without changing it.<br>000b - C0/C1 (No package C-state support)<br>001b - C2<br>010b - C6 (non retention)*<br>011b - C6 (Retention)*<br>100b - Reserved<br>101b - Reserved<br>110b - Reserved<br>111b - No package C-state limit. All C-States supported by the processor are available.<br>Note: C6 retention mode provides more power saving than C6 non-retention mode. Limiting the package to C6 non retention mode does prevent the MSR_PKG_C6_RESIDENCY counter (MSR 3F9h) from being incremented. |
|                  |     | 9:3                        |         | Reserved  |
|                  |     | 10                         |         | I/O MWAIT Redirection Enable (R/W)<br>When set, will map IO_read instructions sent to IO registers at MSR_PMG_IO_CAPTURE_BASE[15:0] to MWAIT instructions.  |
|                  |     | 14:11                      |         | Reserved  |
|                  |     | 15                         |         | CFG Lock (R/O)<br>When set, locks bits [15:0] of this register for further writes until the next reset occurs.  |
|                  |     | 25                         |         | Reserved  |
|                  |     | 26                         |         | C1 State Auto Demotion Enable (R/W)<br>When set, the processor will conditionally demote C3/C6/C7 requests to C1 based on uncore auto-demote information.   |
|                  |     | 27                         |         | Reserved  |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 28                         |        | C1 State Auto Undemotion Enable (R/W)<br>When set, enables Undemotion from Demoted C1.   |
|                  |     | 29                         |        | PKG C-State Auto Demotion Enable (R/W)<br>When set, enables Package C state demotion.  |
|                  |     | 63:30                      |        | Reserved   |
| E4H              | 228 | MSR_PMG_IO_CAPTURE_BASE    | Tile   | Power Management IO Capture Base (R/W)   |
|                  |     | 15:0                       |        | LVL_2 Base Address (R/W)<br>Microcode will compare IO-read zone to this base address to determine if an MWAIT(C2/3/4) needs to be issued instead of the IO-read. Should be programmed to the chipset Plevel_2 IO address.  |
|                  |     | 22:16                      |        | C-State Range (R/W)<br>The IO-port block size in which IO-redirection will be executed (0-127). Should be programmed based on the number of LVLx registers existing in the chipset.  |
|                  |     | 63:23                      |        | Reserved   |
| E7H              | 231 | IA32_MPERF                 | Thread | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.   |
| E8H              | 232 | IA32_APERF                 | Thread | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP               | Core   | Memory Type Range Register (R)<br>See Table 2-2.   |
| 13CH             | 52  | MSR_FEATURE_CONFIG         | Core   | AES Configuration (RW-L)<br>Privileged post-BIOS agent must provide a #GP handler to handle unsuccessful read of this MSR.   |
|                  |     | 1:0                        |        | AES Configuration (RW-L)<br>Upon a successful read of this MSR, the configuration of AES instruction set availability is as follows:<br>11b: AES instructions are not available until next RESET.<br>Otherwise, AES instructions are available.<br>Note, the AES instruction set is not available if read is unsuccessful. If the configuration is not 01b, AES instructions can be mis-configured if a privileged agent unintentionally writes 11b. |
|                  |     | 63:2                       |        | Reserved   |
| 140H             | 320 | MISC_FEATURE_ENABLES       | Thread | MISC_FEATURE_ENABLES   |
|                  |     | 0                          |        | Reserved   |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|-----|----------------------------|--------|---|
| Hex              | Dec |                            |        |   |
|                  |     | 1                          |        | User Mode MONITOR and MWAIT (R/W)<br>If set to 1, the MONITOR and MWAIT instructions do not cause invalid-opcode exceptions when executed with CPL > 0 or in virtual-8086 mode. If MWAIT is executed when CPL > 0 or in virtual-8086 mode, and if EAX indicates a C-state other than C0 or C1, the instruction operates as if EAX indicated the C-state C1. |
|                  |     | 63:2                       |        | Reserved  |
| 174H             | 372 | IA32_SYSENTER_CS           | Thread | See Table 2-2.  |
| 175H             | 373 | IA32_SYSENTER_ESP          | Thread | See Table 2-2.  |
| 176H             | 374 | IA32_SYSENTER_EIP          | Thread | See Table 2-2.  |
| 179H             | 377 | IA32_MCG_CAP               | Thread | See Table 2-2.  |
| 17AH             | 378 | IA32_MCG_STATUS            | Thread | See Table 2-2.  |
| 17DH             | 390 | MSR_SMM_MCA_CAP            | Thread | Enhanced SMM Capabilities (SMM-RO)<br>Reports SMM capability Enhancement. Accessible only while in SMM.   |
|                  |     | 31:0                       |        | Bank Support (SMM-RO)<br>One bit per MCA bank. If the bit is set, that bank supports Enhanced MCA (Default all 0; does not support EMCA).   |
|                  |     | 55:32                      |        | Reserved  |
|                  |     | 56                         |        | Targeted SMI (SMM-RO)<br>Set if targeted SMI is supported.  |
|                  |     | 57                         |        | SMM_CPU_SVRSTR (SMM-RO)<br>Set if SMM SRAM save/restore feature is supported.   |
|                  |     | 58                         |        | SMM_CODE_ACCESS_CHK (SMM-RO)<br>Set if SMM code access check feature is supported.  |
|                  |     | 59                         |        | Long_Flow_Indication (SMM-RO)<br>If set to 1, indicates that the SMM long flow indicator is supported and a host-space interface available to SMM handler.  |
|                  |     | 63:60                      |        | Reserved  |
| 186H             | 390 | IA32_PERFEVTSELO           | Thread | Performance Monitoring Event Select Register (R/W)<br>See Table 2-2.  |
|                  |     | 7:0                        |        | Event Select  |
|                  |     | 15:8                       |        | UMask   |
|                  |     | 16                         |        | USR   |
|                  |     | 17                         |        | OS  |
|                  |     | 18                         |        | Edge  |
|                  |     | 19                         |        | PC  |



**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description  |
|------------------|-----|----------------------------|---------|--|
| Hex              | Dec |                            |         |  |
|                  |     | 20                         |         | INT  |
|                  |     | 21                         |         | AnyThread  |
|                  |     | 22                         |         | EN   |
|                  |     | 23                         |         | INV  |
|                  |     | 31:24                      |         | CMASK  |
|                  |     | 63:32                      |         | Reserved   |
| 187H             | 391 | IA32_PERFEVTSEL1           | Thread  | See Table 2-2.   |
| 198H             | 408 | IA32_PERF_STATUS           | Package | See Table 2-2.   |
| 199H             | 409 | IA32_PERF_CTL              | Thread  | See Table 2-2.   |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Thread  | Clock Modulation (R/W)<br>See Table 2-2.   |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Module  | Thermal Interrupt Control (R/W)<br>See Table 2-2.  |
| 19CH             | 412 | IA32_THERM_STATUS          | Module  | Thermal Monitor Status (R/W)<br>See Table 2-2.   |
|                  |     | 0                          |         | Thermal Status (RO)  |
|                  |     | 1                          |         | Thermal Status Log (R/WCO)   |
|                  |     | 2                          |         | PROTCHOT # or FORCEPR# Status (RO)   |
|                  |     | 3                          |         | PROTCHOT # or FORCEPR# Log (R/WCO)   |
|                  |     | 4                          |         | Critical Temperature Status (RO)   |
|                  |     | 5                          |         | Critical Temperature Status Log (R/WCO)  |
|                  |     | 6                          |         | Thermal Threshold #1 Status (RO)   |
|                  |     | 7                          |         | Thermal Threshold #1 Log (R/WCO)   |
|                  |     | 8                          |         | Thermal Threshold #2 Status (RO)   |
|                  |     | 9                          |         | Thermal Threshold #2 Log (R/WCO)   |
|                  |     | 10                         |         | Power Limitation Status (RO)   |
|                  |     | 11                         |         | Power Limitation Log (R/WCO)   |
|                  |     | 15:12                      |         | Reserved   |
|                  |     | 22:16                      |         | Digital Readout (RO)   |
|                  |     | 26:23                      |         | Reserved   |
|                  |     | 30:27                      |         | Resolution in Degrees Celsius (RO)   |
| 31               |     | Reading Valid (RO)         |         |  |
| 63:32            |     | Reserved                   |         |  |
| 1A0H             | 416 | IA32_MISC_ENABLE           | Thread  | Enable Misc. Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled. |
|                  |     | 0                          |         | Fast-Strings Enable  |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 2:1                        |         | Reserved  |
|                  |     | 3                          |         | Automatic Thermal Control Circuit Enable (R/W)  |
|                  |     | 6:4                        |         | Reserved  |
|                  |     | 7                          |         | Performance Monitoring Available (R)  |
|                  |     | 10:8                       |         | Reserved  |
|                  |     | 11                         |         | Branch Trace Storage Unavailable (RO)   |
|                  |     | 12                         |         | Processor Event Based Sampling Unavailable (RO)   |
|                  |     | 15:13                      |         | Reserved  |
|                  |     | 16                         |         | Enhanced Intel SpeedStep Technology Enable (R/W)  |
|                  |     | 18                         |         | ENABLE MONITOR FSM (R/W)  |
|                  |     | 21:19                      |         | Reserved  |
|                  |     | 22                         |         | Limit CPUID Maxval (R/W)  |
|                  |     | 23                         |         | xTPR Message Disable (R/W)  |
|                  |     | 33:24                      |         | Reserved  |
|                  |     | 34                         |         | XD Bit Disable (R/W)  |
|                  |     | 37:35                      |         | Reserved  |
|                  |     | 38                         |         | Turbo Mode Disable (R/W)  |
|                  |     | 63:39                      |         | Reserved  |
| 1A2H             | 418 | MSR_TEMPERATURE_TARGET     | Package | Temperature Target  |
|                  |     | 15:0                       |         | Reserved  |
|                  |     | 23:16                      |         | Temperature Target (R)  |
|                  |     | 29:24                      |         | Target Offset (R/W)   |
|                  |     | 63:30                      |         | Reserved  |
| 1A4H             | 420 | MSR_MISC_FEATURE_CONTROL   |         | Miscellaneous Feature Control (R/W)   |
|                  |     | 0                          | Core    | DCU Hardware Prefetcher Disable (R/W)<br>If 1, disables the L1 data cache prefetcher.   |
|                  |     | 1                          | Core    | L2 Hardware Prefetcher Disable (R/W)<br>If 1, disables the L2 hardware prefetcher.  |
|                  |     | 63:2                       |         | Reserved  |
| 1A6H             | 422 | MSR_OFFCORE_RSP_0          | Shared  | Offcore Response Event Select Register (R/W)  |
| 1A7H             | 423 | MSR_OFFCORE_RSP_1          | Shared  | Offcore Response Event Select Register (R/W)  |
| 1ADH             | 429 | MSR_TURBO_RATIO_LIMIT      | Package | Maximum Ratio Limit of Turbo Mode for Groups of Cores (RW)  |
|                  |     | 0                          |         | Reserved  |
|                  |     | 7:1                        | Package | Maximum Number of Cores in Group 0<br>Number active processor cores which operates under the maximum ratio limit for group 0. |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|-----|----------------------------|---------|---|
| Hex              | Dec |                            |         |   |
|                  |     | 15:8                       | Package | Maximum Ratio Limit for Group 0<br>Maximum turbo ratio limit when the number of active cores are not more than the group 0 maximum core count.  |
|                  |     | 20:16                      | Package | Number of Incremental Cores Added to Group 1<br>Group 1, which includes the specified number of additional cores plus the cores in group 0, operates under the group 1 turbo max ratio limit = "group 0 Max ratio limit" - "group ratio delta for group 1".     |
|                  |     | 23:21                      | Package | Group Ratio Delta for Group 1<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit to Group 0.   |
|                  |     | 28:24                      | Package | Number of Incremental Cores Added to Group 2<br>Group 2, which includes the specified number of additional cores plus all the cores in group 1, operates under the group 2 turbo max ratio limit = "group 1 Max ratio limit" - "group ratio delta for group 2". |
|                  |     | 31:29                      | Package | Group Ratio Delta for Group 2<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit for Group 1.  |
|                  |     | 36:32                      | Package | Number of Incremental Cores Added to Group 3<br>Group 3, which includes the specified number of additional cores plus all the cores in group 2, operates under the group 3 turbo max ratio limit = "group 2 Max ratio limit" - "group ratio delta for group 3". |
|                  |     | 39:37                      | Package | Group Ratio Delta for Group 3<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit for Group 2.  |
|                  |     | 44:40                      | Package | Number of Incremental Cores Added to Group 4<br>Group 4, which includes the specified number of additional cores plus all the cores in group 3, operates under the group 4 turbo max ratio limit = "group 3 Max ratio limit" - "group ratio delta for group 4". |
|                  |     | 47:45                      | Package | Group Ratio Delta for Group 4<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit for Group 3.  |
|                  |     | 52:48                      | Package | Number of Incremental Cores Added to Group 5<br>Group 5, which includes the specified number of additional cores plus all the cores in group 4, operates under the group 5 turbo max ratio limit = "group 4 Max ratio limit" - "group ratio delta for group 5". |
|                  |     | 55:53                      | Package | Group Ratio Delta for Group 5<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit for Group 4.  |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields   | Scope   | Bit Description   |
|------------------|------|------------------------------|---------|---|
| Hex              | Dec  |                              |         |   |
|                  |      | 60:56                        | Package | Number of Incremental Cores Added to Group 6<br>Group 6, which includes the specified number of additional cores plus all the cores in group 5, operates under the group 6 turbo max ratio limit = "group 5 Max ratio limit" - "group ratio delta for group 6". |
|                  |      | 63:61                        | Package | Group Ratio Delta for Group 6<br>An unsigned integer specifying the ratio decrement relative to the Max ratio limit for Group 5.  |
| 1B0H             | 432  | IA32_ENERGY_PERF_BIAS        | Thread  | See Table 2-2.  |
| 1B1H             | 433  | IA32_PACKAGE_THERM_STATUS    | Package | See Table 2-2.  |
| 1B2H             | 434  | IA32_PACKAGE_THERM_INTERRUPT | Package | See Table 2-2.  |
| 1C8H             | 456  | MSR_LBR_SELECT               | Thread  | Last Branch Record Filtering Select Register (R/W)<br>See Section 17.9.2, "Filtering of Last Branch Records."   |
|                  |      | 0                            |         | CPL_EQ_0  |
|                  |      | 1                            |         | CPL_NEQ_0   |
|                  |      | 2                            |         | JCC   |
|                  |      | 3                            |         | NEAR_REL_CALL   |
|                  |      | 4                            |         | NEAR_IND_CALL   |
|                  |      | 5                            |         | NEAR_RET  |
|                  |      | 6                            |         | NEAR_IND_JMP  |
|                  |      | 7                            |         | NEAR_REL_JMP  |
|                  |      | 8                            |         | FAR_BRANCH  |
|                  | 63:9 |                              |         | Reserved  |
| 1C9H             | 457  | MSR_LASTBRANCH_TOS           | Thread  | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-2) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_0_FROM_IP.  |
| 1D9H             | 473  | IA32_DEBUGCTL                | Thread  | Debug Control (R/W)   |
|                  |      | 0                            |         | LBR<br>Setting this bit to 1 enables the processor to record a running trace of the most recent branches taken by the processor in the LBR stack.   |
|                  |      | 1                            |         | BTF<br>Setting this bit to 1 enables the processor to treat EFLAGS.TF as single-step on branches instead of single-step on instructions.  |
|                  |      | 5:2                          |         | Reserved  |
|                  | 6    |                              |         | TR<br>Setting this bit to 1 enables branch trace messages to be sent.   |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |     | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|-----|----------------------------|--------|--|
| Hex              | Dec |                            |        |  |
|                  |     | 7                          |        | BTS<br>Setting this bit enables branch trace messages (BTMs) to be logged in a BTS buffer.   |
|                  |     | 8                          |        | BTINT<br>When clear, BTMs are logged in a BTS buffer in circular fashion. When this bit is set, an interrupt is generated by the BTS facility when the BTS buffer is full. |
|                  |     | 9                          |        | BTS_OFF_OS<br>When set, BTS or BTM is skipped if CPL = 0.  |
|                  |     | 10                         |        | BTS_OFF_USR<br>When set, BTS or BTM is skipped if CPL > 0.   |
|                  |     | 11                         |        | FREEZE_LBRS_ON_PMI<br>When set, the LBR stack is frozen on a PMI request.  |
|                  |     | 12                         |        | FREEZE_PERFMON_ON_PMI<br>When set, each ENABLE bit of the global counter control MSR are frozen (address 3BFH) on a PMI request.   |
|                  |     | 13                         |        | Reserved   |
|                  |     | 14                         |        | FREEZE_WHILE_SMM<br>When set, freezes perfmon and trace messages while in SMM.   |
|                  |     | 31:15                      |        | Reserved   |
| 1DDH             | 477 | MSR_LER_FROM_LIP           | Thread | Last Exception Record from Linear IP (R)   |
| 1DEH             | 478 | MSR_LER_TO_LIP             | Thread | Last Exception Record to Linear IP (R)   |
| 1F2H             | 498 | IA32_SMRR_PHYSBASE         | Core   | See Table 2-2.   |
| 1F3H             | 499 | IA32_SMRR_PHYSMASK         | Core   | See Table 2-2.   |
| 200H             | 512 | IA32_MTRR_PHYSBASE0        | Core   | See Table 2-2.   |
| 201H             | 513 | IA32_MTRR_PHYSMASK0        | Core   | See Table 2-2.   |
| 202H             | 514 | IA32_MTRR_PHYSBASE1        | Core   | See Table 2-2.   |
| 203H             | 515 | IA32_MTRR_PHYSMASK1        | Core   | See Table 2-2.   |
| 204H             | 516 | IA32_MTRR_PHYSBASE2        | Core   | See Table 2-2.   |
| 205H             | 517 | IA32_MTRR_PHYSMASK2        | Core   | See Table 2-2.   |
| 206H             | 518 | IA32_MTRR_PHYSBASE3        | Core   | See Table 2-2.   |
| 207H             | 519 | IA32_MTRR_PHYSMASK3        | Core   | See Table 2-2.   |
| 208H             | 520 | IA32_MTRR_PHYSBASE4        | Core   | See Table 2-2.   |
| 209H             | 521 | IA32_MTRR_PHYSMASK4        | Core   | See Table 2-2.   |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5        | Core   | See Table 2-2.   |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5        | Core   | See Table 2-2.   |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6        | Core   | See Table 2-2.   |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
| 20DH             | 525  | IA32_MTRR_PHYSMASK6        | Core    | See Table 2-2.  |
| 20EH             | 526  | IA32_MTRR_PHYSBASE7        | Core    | See Table 2-2.  |
| 20FH             | 527  | IA32_MTRR_PHYSMASK7        | Core    | See Table 2-2.  |
| 250H             | 592  | IA32_MTRR_FIX64K_00000     | Core    | See Table 2-2.  |
| 258H             | 600  | IA32_MTRR_FIX16K_80000     | Core    | See Table 2-2.  |
| 259H             | 601  | IA32_MTRR_FIX16K_A0000     | Core    | See Table 2-2.  |
| 268H             | 616  | IA32_MTRR_FIX4K_C0000      | Core    | See Table 2-2.  |
| 269H             | 617  | IA32_MTRR_FIX4K_C8000      | Core    | See Table 2-2.  |
| 26AH             | 618  | IA32_MTRR_FIX4K_D0000      | Core    | See Table 2-2.  |
| 26BH             | 619  | IA32_MTRR_FIX4K_D8000      | Core    | See Table 2-2.  |
| 26CH             | 620  | IA32_MTRR_FIX4K_E0000      | Core    | See Table 2-2.  |
| 26DH             | 621  | IA32_MTRR_FIX4K_E8000      | Core    | See Table 2-2.  |
| 26EH             | 622  | IA32_MTRR_FIX4K_F0000      | Core    | See Table 2-2.  |
| 26FH             | 623  | IA32_MTRR_FIX4K_F8000      | Core    | See Table 2-2.  |
| 277H             | 631  | IA32_PAT                   | Core    | See Table 2-2.  |
| 2FFH             | 767  | IA32_MTRR_DEF_TYPE         | Core    | Default Memory Types (R/W)<br>See Table 2-2.  |
| 309H             | 777  | IA32_FIXED_CTR0            | Thread  | Fixed-Function Performance Counter Register 0 (R/W)<br>See Table 2-2.   |
| 30AH             | 778  | IA32_FIXED_CTR1            | Thread  | Fixed-Function Performance Counter Register 1 (R/W)<br>See Table 2-2.   |
| 30BH             | 779  | IA32_FIXED_CTR2            | Thread  | Fixed-Function Performance Counter Register 2 (R/W)<br>See Table 2-2.   |
| 345H             | 837  | IA32_PERF_CAPABILITIES     | Package | See Table 2-2. See Section 17.4.1, "IA32_DEBUGCTL MSR."   |
| 38DH             | 909  | IA32_FIXED_CTR_CTRL        | Thread  | Fixed-Function-Counter Control Register (R/W)<br>See Table 2-2.   |
| 38EH             | 910  | IA32_PERF_GLOBAL_STATUS    | Thread  | See Table 2-2.  |
| 38FH             | 911  | IA32_PERF_GLOBAL_CTRL      | Thread  | See Table 2-2.  |
| 390H             | 912  | IA32_PERF_GLOBAL_OVF_CTRL  | Thread  | See Table 2-2.  |
| 3F1H             | 1009 | MSR_PEBBS_ENABLE           | Thread  | See Table 2-2.  |
| 3F8H             | 1016 | MSR_PKG_C3_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states. |
|                  |      | 63:0                       |         | Package C3 Residency Counter (R/O)  |
| 3F9H             | 1017 | MSR_PKG_C6_RESIDENCY       | Package |   |
|                  |      | 63:0                       |         | Package C6 Residency Counter (R/O)  |
| 3FAH             | 1018 | MSR_PKG_C7_RESIDENCY       | Package |   |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 63:0                       |         | Package C7 Residency Counter (R/O)  |
| 3FCH             | 1020 | MSR_MC0_RESIDENCY          | Module  | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |
|                  |      | 63:0                       |         | Module C0 Residency Counter (R/O)   |
| 3FDH             | 1021 | MSR_MC6_RESIDENCY          | Module  |   |
|                  |      | 63:0                       |         | Module C6 Residency Counter (R/O)   |
| 3FFH             | 1023 | MSR_CORE_C6_RESIDENCY      | Core    | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |
|                  |      | 63:0                       |         | CORE C6 Residency Counter (R/O)   |
| 400H             | 1024 | IA32_MC0_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 401H             | 1025 | IA32_MC0_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 402H             | 1026 | IA32_MC0_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 404H             | 1028 | IA32_MC1_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 408H             | 1032 | IA32_MC2_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 409H             | 1033 | IA32_MC2_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40AH             | 1034 | IA32_MC2_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 40CH             | 1036 | IA32_MC3_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 40DH             | 1037 | IA32_MC3_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 40EH             | 1038 | IA32_MC3_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 410H             | 1040 | IA32_MC4_CTL               | Core    | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC4_STATUS            | Core    | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC4_ADDR              | Core    | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 414H             | 1044 | IA32_MC5_CTL               | Package | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 415H             | 1045 | IA32_MC5_STATUS            | Package | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 416H             | 1046 | IA32_MC5_ADDR              | Package | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."   |
| 4C1H             | 1217 | IA32_A_PMC0                | Thread  | See Table 2-2.  |
| 4C2H             | 1218 | IA32_A_PMC1                | Thread  | See Table 2-2.  |
| 600H             | 1536 | IA32_DS_AREA               | Thread  | DS Save Area (R/W)<br>See Table 2-2.  |
| 606H             | 1542 | MSR_RAPL_POWER_UNIT        | Package | Unit Multipliers Used in RAPL Interfaces (R/O)  |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope   | Bit Description   |
|------------------|------|----------------------------|---------|---|
| Hex              | Dec  |                            |         |   |
|                  |      | 3:0                        | Package | Power Units<br>See Section 14.9.1, "RAPL Interfaces."   |
|                  |      | 7:4                        | Package | Reserved  |
|                  |      | 12:8                       | Package | Energy Status Units<br>Energy related information (in Joules) is based on the multiplier, $1/2^{\wedge}ESU$ ; where ESU is an unsigned integer represented by bits 12:8. Default value is 0EH (or 61 micro-joules).   |
|                  |      | 15:13                      | Package | Reserved  |
|                  |      | 19:16                      | Package | Time Units<br>See Section 14.9.1, "RAPL Interfaces."  |
|                  |      | 63:20                      |         | Reserved  |
| 60DH             | 1549 | MSR_PKG_C2_RESIDENCY       | Package | Note: C-state values are processor specific C-state code names, unrelated to MWAIT extension C-state parameters or ACPI C-states.   |
|                  |      | 63:0                       |         | Package C2 Residency Counter (R/O)  |
| 610H             | 1552 | MSR_PKG_POWER_LIMIT        | Package | PKG RAPL Power Limit Control (R/W)<br>See Section 14.9.3, "Package RAPL Domain."  |
| 611H             | 1553 | MSR_PKG_ENERGY_STATUS      | Package | PKG Energy Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 613H             | 1555 | MSR_PKG_PERF_STATUS        | Package | PKG Perf Status (R/O)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 614H             | 1556 | MSR_PKG_POWER_INFO         | Package | PKG RAPL Parameters (R/W)<br>See Section 14.9.3, "Package RAPL Domain."   |
| 618H             | 1560 | MSR_DRAM_POWER_LIMIT       | Package | DRAM RAPL Power Limit Control (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."  |
| 619H             | 1561 | MSR_DRAM_ENERGY_STATUS     | Package | DRAM Energy Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61BH             | 1563 | MSR_DRAM_PERF_STATUS       | Package | DRAM Performance Throttling Status (R/O)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 61CH             | 1564 | MSR_DRAM_POWER_INFO        | Package | DRAM RAPL Parameters (R/W)<br>See Section 14.9.5, "DRAM RAPL Domain."   |
| 620H             | 1568 | MSR_UNCORE_RATIO_LIMIT     | Package | Uncore Ratio Limit (R/W)<br>Out of reset, the min_ratio and max_ratio fields represent the widest possible range of uncore frequencies. Writing to these fields allows software to control the minimum and the maximum frequency that hardware will select. |
|                  |      | 63:15                      |         | Reserved  |



**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields  | Scope   | Bit Description   |
|------------------|------|-----------------------------|---------|---|
| Hex              | Dec  |                             |         |   |
|                  |      | 14:8                        |         | MIN_RATIO<br>Writing to this field controls the minimum possible ratio of the LLC/Ring.                     |
|                  |      | 7                           |         | Reserved  |
|                  |      | 6:0                         |         | MAX_RATIO<br>This field is used to limit the max ratio of the LLC/Ring.                                     |
| 638H             | 1592 | MSR_PPO_POWER_LIMIT         | Package | PPO RAPL Power Limit Control (R/W)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."                           |
| 639H             | 1593 | MSR_PPO_ENERGY_STATUS       | Package | PPO Energy Status (R/O)<br>See Section 14.9.4, "PPO/PP1 RAPL Domains."                                      |
| 648H             | 1608 | MSR_CONFIG_TDP_NOMINAL      | Package | Base TDP Ratio (R/O)<br>See Table 2-24.   |
| 649H             | 1609 | MSR_CONFIG_TDP_LEVEL1       | Package | ConfigTDP Level 1 ratio and power level (R/O)<br>See Table 2-24.  |
| 64AH             | 1610 | MSR_CONFIG_TDP_LEVEL2       | Package | ConfigTDP Level 2 ratio and power level (R/O)<br>See Table 2-24.  |
| 64BH             | 1611 | MSR_CONFIG_TDP_CONTROL      | Package | ConfigTDP Control (R/W)<br>See Table 2-24.  |
| 64CH             | 1612 | MSR_TURBO_ACTIVATION_RATIO  | Package | ConfigTDP Control (R/W)<br>See Table 2-24.  |
| 690H             | 1680 | MSR_CORE_PERF_LIMIT_REASONS | Package | Indicator of Frequency Clipping in Processor Cores (R/W)<br>(Frequency refers to processor core frequency.) |
|                  |      | 0                           |         | PROCHOT Status (R0)   |
|                  |      | 1                           |         | Thermal Status (R0)   |
|                  |      | 5:2                         |         | Reserved  |
|                  |      | 6                           |         | VR Therm Alert Status (R0)  |
|                  |      | 7                           |         | Reserved  |
|                  |      | 8                           |         | Electrical Design Point Status (R0)   |
|                  |      | 63:9                        |         | Reserved  |
| 6E0H             | 1760 | IA32_TSC_DEADLINE           | Core    | TSC Target of Local APIC's TSC Deadline Mode (R/W)<br>See Table 2-2.  |
| 802H             | 2050 | IA32_X2APIC_APICID          | Thread  | x2APIC ID Register (R/O)  |
| 803H             | 2051 | IA32_X2APIC_VERSION         | Thread  | x2APIC Version Register (R/O)   |
| 808H             | 2056 | IA32_X2APIC_TPR             | Thread  | x2APIC Task Priority Register (R/W)   |
| 80AH             | 2058 | IA32_X2APIC_PPR             | Thread  | x2APIC Processor Priority Register (R/O)  |
| 80BH             | 2059 | IA32_X2APIC_EOI             | Thread  | x2APIC EOI Register (W/O)   |
| 80DH             | 2061 | IA32_X2APIC_LDR             | Thread  | x2APIC Logical Destination Register (R/O)   |
| 80FH             | 2063 | IA32_X2APIC_SIVR            | Thread  | x2APIC Spurious Interrupt Vector Register (R/W)   |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description   |
|------------------|------|----------------------------|--------|---|
| Hex              | Dec  |                            |        |   |
| 810H             | 2064 | IA32_X2APIC_ISR0           | Thread | x2APIC In-Service Register Bits [31:0] (R/O)                |
| 811H             | 2065 | IA32_X2APIC_ISR1           | Thread | x2APIC In-Service Register Bits [63:32] (R/O)               |
| 812H             | 2066 | IA32_X2APIC_ISR2           | Thread | x2APIC In-Service Register Bits [95:64] (R/O)               |
| 813H             | 2067 | IA32_X2APIC_ISR3           | Thread | x2APIC In-Service Register Bits [127:96] (R/O)              |
| 814H             | 2068 | IA32_X2APIC_ISR4           | Thread | x2APIC In-Service Register Bits [159:128] (R/O)             |
| 815H             | 2069 | IA32_X2APIC_ISR5           | Thread | x2APIC In-Service Register Bits [191:160] (R/O)             |
| 816H             | 2070 | IA32_X2APIC_ISR6           | Thread | x2APIC In-Service Register Bits [223:192] (R/O)             |
| 817H             | 2071 | IA32_X2APIC_ISR7           | Thread | x2APIC In-Service Register Bits [255:224] (R/O)             |
| 818H             | 2072 | IA32_X2APIC_TMR0           | Thread | x2APIC Trigger Mode Register Bits [31:0] (R/O)              |
| 819H             | 2073 | IA32_X2APIC_TMR1           | Thread | x2APIC Trigger Mode Register Bits [63:32] (R/O)             |
| 81AH             | 2074 | IA32_X2APIC_TMR2           | Thread | x2APIC Trigger Mode Register Bits [95:64] (R/O)             |
| 81BH             | 2075 | IA32_X2APIC_TMR3           | Thread | x2APIC Trigger Mode Register Bits [127:96] (R/O)            |
| 81CH             | 2076 | IA32_X2APIC_TMR4           | Thread | x2APIC Trigger Mode Register Bits [159:128] (R/O)           |
| 81DH             | 2077 | IA32_X2APIC_TMR5           | Thread | x2APIC Trigger Mode Register Bits [191:160] (R/O)           |
| 81EH             | 2078 | IA32_X2APIC_TMR6           | Thread | x2APIC Trigger Mode Register Bits [223:192] (R/O)           |
| 81FH             | 2079 | IA32_X2APIC_TMR7           | Thread | x2APIC Trigger Mode Register Bits [255:224] (R/O)           |
| 820H             | 2080 | IA32_X2APIC_IRR0           | Thread | x2APIC Interrupt Request Register Bits [31:0] (R/O)         |
| 821H             | 2081 | IA32_X2APIC_IRR1           | Thread | x2APIC Interrupt Request Register Bits [63:32] (R/O)        |
| 822H             | 2082 | IA32_X2APIC_IRR2           | Thread | x2APIC Interrupt Request Register Bits [95:64] (R/O)        |
| 823H             | 2083 | IA32_X2APIC_IRR3           | Thread | x2APIC Interrupt Request Register Bits [127:96] (R/O)       |
| 824H             | 2084 | IA32_X2APIC_IRR4           | Thread | x2APIC Interrupt Request Register Bits [159:128] (R/O)      |
| 825H             | 2085 | IA32_X2APIC_IRR5           | Thread | x2APIC Interrupt Request Register Bits [191:160] (R/O)      |
| 826H             | 2086 | IA32_X2APIC_IRR6           | Thread | x2APIC Interrupt Request Register Bits [223:192] (R/O)      |
| 827H             | 2087 | IA32_X2APIC_IRR7           | Thread | x2APIC Interrupt Request Register Bits [255:224] (R/O)      |
| 828H             | 2088 | IA32_X2APIC_ESR            | Thread | x2APIC Error Status Register (R/W)                          |
| 82FH             | 2095 | IA32_X2APIC_LVT_CMCI       | Thread | x2APIC LVT Corrected Machine Check Interrupt Register (R/W) |
| 830H             | 2096 | IA32_X2APIC_ICR            | Thread | x2APIC Interrupt Command Register (R/W)                     |
| 832H             | 2098 | IA32_X2APIC_LVT_TIMER      | Thread | x2APIC LVT Timer Interrupt Register (R/W)                   |
| 833H             | 2099 | IA32_X2APIC_LVT_THERMAL    | Thread | x2APIC LVT Thermal Sensor Interrupt Register (R/W)          |
| 834H             | 2100 | IA32_X2APIC_LVT_PMI        | Thread | x2APIC LVT Performance Monitor Register (R/W)               |
| 835H             | 2101 | IA32_X2APIC_LVT_LINT0      | Thread | x2APIC LVT LINT0 Register (R/W)                             |
| 836H             | 2102 | IA32_X2APIC_LVT_LINT1      | Thread | x2APIC LVT LINT1 Register (R/W)                             |
| 837H             | 2103 | IA32_X2APIC_LVT_ERROR      | Thread | x2APIC LVT Error Register (R/W)                             |
| 838H             | 2104 | IA32_X2APIC_INIT_COUNT     | Thread | x2APIC Initial Count Register (R/W)                         |
| 839H             | 2105 | IA32_X2APIC_CUR_COUNT      | Thread | x2APIC Current Count Register (R/O)                         |

**Table 2-43. Selected MSRs Supported by Intel® Xeon Phi™ Processors with DisplayFamily\_DisplayModel Signatures 06\_57H and 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope  | Bit Description  |
|------------------|------|----------------------------|--------|--|
| Hex              | Dec  |                            |        |  |
| 83EH             | 2110 | IA32_X2APIC_DIV_CONF       | Thread | x2APIC Divide Configuration Register (R/W)                     |
| 83FH             | 2111 | IA32_X2APIC_SELF_IPI       | Thread | x2APIC Self IPI Register (W/O)                                 |
| C000_0080H       |      | IA32_EFER                  | Thread | Extended Feature Enables<br>See Table 2-2.                     |
| C000_0081H       |      | IA32_STAR                  | Thread | System Call Target Address (R/W)<br>See Table 2-2.             |
| C000_0082H       |      | IA32_LSTAR                 | Thread | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2. |
| C000_0084H       |      | IA32_FMASK                 | Thread | System Call Flag Mask (R/W)<br>See Table 2-2.                  |
| C000_0100H       |      | IA32_FS_BASE               | Thread | Map of BASE Address of FS (R/W)<br>See Table 2-2.              |
| C000_0101H       |      | IA32_GS_BASE               | Thread | Map of BASE Address of GS (R/W)<br>See Table 2-2.              |
| C000_0102H       |      | IA32_KERNEL_GS_BASE        | Thread | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.      |
| C000_0103H       |      | IA32_TSC_AUX               | Thread | AUXILIARY TSC Signature (R/W)<br>See Table 2-2                 |

Table 2-44 lists model-specific registers that are supported by Intel® Xeon Phi™ processor 7215, 7285, 7295 series based on the Knights Mill microarchitecture.

**Table 2-44. Additional MSRs Supported by Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series with DisplayFamily\_DisplayModel Signature 06\_85H**

| Register Address |      | Register Name / Bit Fields | Scope | Bit Description  |
|------------------|------|----------------------------|-------|--|
| Hex              | Dec  |                            |       |  |
| 9BH              | 155  | IA32_SMM_MONITOR_CTL       | Core  | SMM Monitor Configuration (R/W)<br>This MSR is readable only if VMX is enabled, and writeable only if VMX is enabled and in SMM mode, and is used to configure the VMX MSEG base address. See Table 2-2. |
| 480H             | 1152 | IA32_VMX_BASIC             | Core  | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.   |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL      | Core  | Capability Reporting Register of Pin-based VM-execution Controls (R/O)<br>See Table 2-2.   |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL     | Core  | Capability Reporting Register of Primary Processor-based VM-execution Controls (R/O)   |
| 483H             | 1155 | IA32_VMX_EXIT_CTL          | Core  | Capability Reporting Register of VM-exit Controls (R/O)<br>See Table 2-2.  |

**Table 2-44. Additional MSRs Supported by Intel® Xeon Phi™ Processor 7215, 7285, 7295 Series with DisplayFamily\_DisplayModel Signature 06\_85H**

| Register Address |      | Register Name / Bit Fields   | Scope | Bit Description   |
|------------------|------|------------------------------|-------|---|
| Hex              | Dec  |                              |       |   |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL5          | Core  | Capability Reporting Register of VM-entry Controls (R/O)<br>See Table 2-2.                                  |
| 485H             | 1157 | IA32_VMX_MISC                | Core  | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Table 2-2.                                |
| 486H             | 1158 | IA32_VMX_CRO_FIXED0          | Core  | Capability Reporting Register of CRO Bits Fixed to 0 (R/O)<br>See Table 2-2.                                |
| 487H             | 1159 | IA32_VMX_CRO_FIXED1          | Core  | Capability Reporting Register of CRO Bits Fixed to 1 (R/O)<br>See Table 2-2.                                |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0          | Core  | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Table 2-2.                                |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1          | Core  | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Table 2-2.                                |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM           | Core  | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Table 2-2.                             |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL52     | Core  | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Table 2-2.    |
| 48CH             | 1164 | IA32_VMX_EPT_VPID_ENUM       | Core  | Capability Reporting Register of EPT and VPID (R/O)<br>See Table 2-2.                                       |
| 48DH             | 1165 | IA32_VMX_TRUE_PINBASE D_CTL5 | Core  | Capability Reporting Register of Pin-Based VM-Execution Flex Controls (R/O)<br>See Table 2-2.               |
| 48EH             | 1166 | IA32_VMX_TRUE_PROCBASED_CTL5 | Core  | Capability Reporting Register of Primary Processor-Based VM-Execution Flex Controls (R/O)<br>See Table 2-2. |
| 48FH             | 1167 | IA32_VMX_TRUE_EXIT_CTL5      | Core  | Capability Reporting Register of VM-Exit Flex Controls (R/O)<br>See Table 2-2.                              |
| 490H             | 1168 | IA32_VMX_TRUE_ENTRY_CTL5     | Core  | Capability Reporting Register of VM-Entry Flex Controls (R/O)<br>See Table 2-2.                             |
| 491H             | 1169 | IA32_VMX_FMFUNC              | Core  | Capability Reporting Register of VM-Function Controls (R/O)<br>See Table 2-2.                               |

## 2.18 MSRS IN THE PENTIUM® 4 AND INTEL® XEON® PROCESSORS

Table 2-45 lists MSRs (architectural and model-specific) that are defined across processor generations based on Intel NetBurst microarchitecture. The processor can be identified by its CPUID signatures of DisplayFamily encoding of 0FH, see Table 2-1.

- MSRs with an “IA32\_” prefix are designated as “architectural.” This means that the functions of these MSRs and their addresses remain the same for succeeding families of IA-32 processors.
- MSRs with an “MSR\_” prefix are model specific with respect to address functionalities. The column “Model Availability” lists the model encoding value(s) within the Pentium 4 and Intel Xeon processor family at the specified register address. The model encoding value of a processor can be queried using CPUID. See “CPUID—CPU Identification” in Chapter 3 of the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2A*.

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors**

| Register Address |     | Register Name<br>Fields and Flags | Model Avail-<br>ability | Shared/<br>Unique <sup>1</sup> | Bit Description  |
|------------------|-----|-----------------------------------|-------------------------|--------------------------------|--|
| Hex              | Dec |                                   |                         |                                |  |
| 0H               | 0   | IA32_P5_MC_ADDR                   | 0, 1, 2, 3,<br>4, 6     | Shared                         | See Section 2.22, “MSRs in Pentium Processors.”  |
| 1H               | 1   | IA32_P5_MC_TYPE                   | 0, 1, 2, 3,<br>4, 6     | Shared                         | See Section 2.22, “MSRs in Pentium Processors.”  |
| 6H               | 6   | IA32_MONITOR_FILTER_LINE_SIZE     | 3, 4, 6                 | Shared                         | See Section 8.10.5, “Monitor/Mwait Address Range Determination.”   |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER           | 0, 1, 2, 3,<br>4, 6     | Unique                         | Time Stamp Counter<br>See Table 2-2.   |
|                  |     |                                   |                         |                                | On earlier processors, only the lower 32 bits are writable. On any write to the lower 32 bits, the upper 32 bits are cleared. For processor family 0FH, models 3 and 4: all 64 bits are writable.  |
| 17H              | 23  | IA32_PLATFORM_ID                  | 0, 1, 2, 3,<br>4, 6     | Shared                         | Platform ID (R)<br>See Table 2-2.<br><br>The operating system can use this MSR to determine “slot” information for the processor and the proper microcode update to load.  |
| 1BH              | 27  | IA32_APIC_BASE                    | 0, 1, 2, 3,<br>4, 6     | Unique                         | APIC Location and Status (R/W)<br>See Table 2-2. See Section 10.4.4, “Local APIC Status and Location.”   |
| 2AH              | 42  | MSR_EBC_HARD_POWERON              | 0, 1, 2, 3,<br>4, 6     | Shared                         | Processor Hard Power-On Configuration (R/W) Enables and disables processor features. (R) Indicates current processor configuration.  |
|                  |     | 0                                 |                         |                                | Output Tri-state Enabled (R)<br>Indicates whether tri-state output is enabled (1) or disabled (0) as set by the strapping of SMI#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted. |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description  |
|------------------|-----|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec |                                |                    |                            |  |
|                  |     | 1                              |                    |                            | Execute BIST (R)<br>Indicates whether the execution of the BIST is enabled (1) or disabled (0) as set by the strapping of INIT#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted.                   |
|                  |     | 2                              |                    |                            | In Order Queue Depth (R)<br>Indicates whether the in order queue depth for the system bus is 1 (1) or up to 12 (0) as set by the strapping of A7#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted. |
|                  |     | 3                              |                    |                            | MCERR# Observation Disabled (R)<br>Indicates whether MCERR# observation is enabled (0) or disabled (1) as determined by the strapping of A9#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted.      |
|                  |     | 4                              |                    |                            | BINIT# Observation Enabled (R)<br>Indicates whether BINIT# observation is enabled (0) or disabled (1) as determined by the strapping of A10#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted.      |
|                  |     | 6:5                            |                    |                            | APIC Cluster ID (R)<br>Contains the logical APIC cluster ID value as set by the strapping of A12# and A11#. The logical cluster ID value is written into the field on the deassertion of RESET#; the field is set to 1 when the address bus signal is asserted.                |
|                  |     | 7                              |                    |                            | Bus Park Disable (R)<br>Indicates whether bus park is enabled (0) or disabled (1) as set by the strapping of A15#. The value in this bit is written on the deassertion of RESET#; the bit is set to 1 when the address bus signal is asserted.                                 |
|                  |     | 11:8                           |                    |                            | Reserved   |
|                  |     | 13:12                          |                    |                            | Agent ID (R)<br>Contains the logical agent ID value as set by the strapping of BR[3:0]. The logical ID value is written into the field on the deassertion of RESET#; the field is set to 1 when the address bus signal is asserted.  |
|                  |     | 63:14                          |                    |                            | Reserved   |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description   |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|---|
| Hex              | Dec |                                   |                            |                                |   |
| 2BH              | 43  | MSR_EBC_SOFT_POWERON              | 0, 1, 2, 3,<br>4, 6        | Shared                         | Processor Soft Power-On Configuration (R/W)<br>Enables and disables processor features.   |
|                  |     | 0                                 |                            |                                | RCNT/SCNT On Request Encoding Enable (R/W)<br>Controls the driving of RCNT/SCNT on the<br>request encoding. Set to enable (1); clear to<br>disabled (0, default).   |
|                  |     | 1                                 |                            |                                | Data Error Checking Disable (R/W)<br>Set to disable system data bus parity checking;<br>clear to enable parity checking.  |
|                  |     | 2                                 |                            |                                | Response Error Checking Disable (R/W)<br>Set to disable (default); clear to enable.   |
|                  |     | 3                                 |                            |                                | Address/Request Error Checking Disable (R/W)<br>Set to disable (default); clear to enable.  |
|                  |     | 4                                 |                            |                                | Initiator MCERR# Disable (R/W)<br>Set to disable MCERR# driving for initiator bus<br>requests (default); clear to enable.   |
|                  |     | 5                                 |                            |                                | Internal MCERR# Disable (R/W)<br>Set to disable MCERR# driving for initiator<br>internal errors (default); clear to enable.   |
|                  |     | 6                                 |                            |                                | BINIT# Driver Disable (R/W)<br>Set to disable BINIT# driver (default); clear to<br>enable driver.   |
|                  |     | 63:7                              |                            |                                | Reserved  |
| 2CH              | 44  | MSR_EBC_FREQUENCY_ID              | 2,3, 4, 6                  | Shared                         | Processor Frequency Configuration<br>The bit field layout of this MSR varies according<br>to the MODEL value in the CPUID version<br>information. The following bit field layout<br>applies to Pentium 4 and Xeon Processors with<br>MODEL encoding equal or greater than 2.<br>(R) The field Indicates the current processor<br>frequency configuration. |
|                  |     | 15:0                              |                            |                                | Reserved  |
|                  |     | 18:16                             |                            |                                | Scalable Bus Speed (R/W)<br>Indicates the intended scalable bus speed:<br><u>Encoding Scalable Bus Speed</u><br>000B 100 MHz (Model 2)<br>000B 266 MHz (Model 3 or 4)<br>001B 133 MHz<br>010B 200 MHz<br>011B 166 MHz<br>100B 333 MHz (Model 6)   |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description   |
|------------------|-----|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec |                                |                    |                            |   |
|                  |     |                                |                    |                            | 133.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 001B.<br>166.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 011B.  |
|                  |     |                                |                    |                            | 266.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 000B and model encoding = 3 or 4.<br>333.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 100B and model encoding = 6.<br>All other values are reserved.       |
|                  |     | 23:19                          |                    |                            | Reserved  |
|                  |     | 31:24                          |                    |                            | Core Clock Frequency to System Bus Frequency Ratio (R)<br>The processor core clock frequency to system bus frequency ratio observed at the de-assertion of the reset pin.   |
|                  |     | 63:25                          |                    |                            | Reserved  |
| 2CH              | 44  | MSR_EBC_FREQUENCY_ID           | 0, 1               | Shared                     | Processor Frequency Configuration (R)<br>The bit field layout of this MSR varies according to the MODEL value of the CPUID version information. This bit field layout applies to Pentium 4 and Xeon Processors with MODEL encoding less than 2.<br>Indicates current processor frequency configuration. |
|                  |     | 20:0                           |                    |                            | Reserved  |
|                  |     | 23:21                          |                    |                            | Scalable Bus Speed (R/W)<br>Indicates the intended scalable bus speed:<br><u>Encoding Scalable Bus Speed</u><br>000B 100 MHz<br><br>All others values reserved.   |
|                  |     | 63:24                          |                    |                            | Reserved  |
| 3AH              | 58  | IA32_FEATURE_CONTROL           | 3, 4, 6            | Unique                     | Control Features in IA-32 Processor (R/W)<br>See Table 2-2.<br>(If CPUID.01H:ECX.[bit 5])   |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG            | 0, 1, 2, 3, 4, 6   | Shared                     | BIOS Update Trigger Register (W)<br>See Table 2-2.  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID              | 0, 1, 2, 3, 4, 6   | Unique                     | BIOS Update Signature ID (R/W)<br>See Table 2-2.  |



Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description  |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|--|
| Hex              | Dec |                                   |                            |                                |  |
| 9BH              | 155 | IA32_SMM_MONITOR_CTL              | 3, 4, 6                    | Unique                         | SMM Monitor Configuration (R/W)<br>See Table 2-2.  |
| FEH              | 254 | IA32_MTRRCAP                      | 0, 1, 2, 3,<br>4, 6        | Unique                         | MTRR Information<br>See Section 11.11.1, "MTRR Feature<br>Identification."   |
| 174H             | 372 | IA32_SYSENTER_CS                  | 0, 1, 2, 3,<br>4, 6        | Unique                         | CS Register Target for CPL 0 Code (R/W)<br>See Table 2-2.<br>See Section 5.8.7, "Performing Fast Calls to<br>System Procedures with the SYSENTER and<br>SYSEXIT Instructions." |
| 175H             | 373 | IA32_SYSENTER_ESP                 | 0, 1, 2, 3,<br>4, 6        | Unique                         | Stack Pointer for CPL 0 Stack (R/W)<br>See Table 2-2.<br>See Section 5.8.7, "Performing Fast Calls to<br>System Procedures with the SYSENTER and<br>SYSEXIT Instructions."     |
| 176H             | 374 | IA32_SYSENTER_EIP                 | 0, 1, 2, 3,<br>4, 6        | Unique                         | CPL 0 Code Entry Point (R/W)<br>See Table 2-2. See Section 5.8.7, "Performing<br>Fast Calls to System Procedures with the<br>SYSENTER and SYSEXIT Instructions."               |
| 179H             | 377 | IA32_MCG_CAP                      | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check Capabilities (R)<br>See Table 2-2. See Section 15.3.1.1,<br>"IA32_MCG_CAP MSR."  |
| 17AH             | 378 | IA32_MCG_STATUS                   | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check Status (R)<br>See Table 2-2. See Section 15.3.1.2,<br>"IA32_MCG_STATUS MSR."   |
| 17BH             | 379 | IA32_MCG_CTL                      |                            |                                | Machine Check Feature Enable (R/W)<br>See Table 2-2.<br>See Section 15.3.1.3, "IA32_MCG_CTL MSR."  |
| 180H             | 384 | MSR_MCG_RAX                       | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check EAX/RAX Save State<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."   |
|                  |     | 63:0                              |                            |                                | Contains register state at time of machine check<br>error. When in non-64-bit modes at the time of<br>the error, bits 63-32 do not contain valid data.                         |
| 181H             | 385 | MSR_MCG_RBX                       | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check EBX/RBX Save State<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."   |
|                  |     | 63:0                              |                            |                                | Contains register state at time of machine check<br>error. When in non-64-bit modes at the time of<br>the error, bits 63-32 do not contain valid data.                         |
| 182H             | 386 | MSR_MCG_RCX                       | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check ECX/RCX Save State<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."   |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description  |
|------------------|-----|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec |                                |                    |                            |  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 183H             | 387 | MSR_MCG_RDX                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check EDX/RDX Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 184H             | 388 | MSR_MCG_RSI                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check ESI/RSI Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 185H             | 389 | MSR_MCG_RDI                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check EDI/RDI Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 186H             | 390 | MSR_MCG_RBP                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check EBP/RBP Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 187H             | 391 | MSR_MCG_RSP                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check ESP/RSP Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 188H             | 392 | MSR_MCG_RFLAGS                 | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check EFLAGS/RFLAG Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."                                     |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |
| 189H             | 393 | MSR_MCG_RIP                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check EIP/RIP Save State<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Contains register state at time of machine check error. When in non-64-bit modes at the time of the error, bits 63-32 do not contain valid data. |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags        | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description  |
|------------------|-----|--|----------------------------|--------------------------------|--|
| Hex              | Dec |  |                            |                                |  |
| 18AH             | 394 | MSR_MCG_MISC                             | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check Miscellaneous<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."  |
|                  |     | 0  |                            |                                | DS<br>When set, the bit indicates that a page assist or<br>page fault occurred during DS normal operation.<br>The processors response is to shut down.<br>The bit is used as an aid for debugging DS<br>handling code. It is the responsibility of the user<br>(BIOS or operating system) to clear this bit for<br>normal operation. |
|                  |     | 63:1                                     |                            |                                | Reserved   |
| 18BH-<br>18FH    | 395 | MSR_MCG_RESERVED1 -<br>MSR_MCG_RESERVED5 |                            |                                | Reserved   |
| 190H             | 400 | MSR_MCG_R8                               | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check R8<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."   |
|                  |     | 63:0                                     |                            |                                | Registers R8-15 (and the associated state-save<br>MSRs) exist only in Intel 64 processors. These<br>registers contain valid information only when<br>the processor is operating in 64-bit mode at the<br>time of the error.  |
| 191H             | 401 | MSR_MCG_R9                               | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check R9D/R9<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."   |
|                  |     | 63:0                                     |                            |                                | Registers R8-15 (and the associated state-save<br>MSRs) exist only in Intel 64 processors. These<br>registers contain valid information only when<br>the processor is operating in 64-bit mode at the<br>time of the error.  |
| 192H             | 402 | MSR_MCG_R10                              | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check R10<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."  |
|                  |     | 63:0                                     |                            |                                | Registers R8-15 (and the associated state-save<br>MSRs) exist only in Intel 64 processors. These<br>registers contain valid information only when<br>the processor is operating in 64-bit mode at the<br>time of the error.  |
| 193H             | 403 | MSR_MCG_R11                              | 0, 1, 2, 3,<br>4, 6        | Unique                         | Machine Check R11<br>See Section 15.3.2.6, "IA32_MCG Extended<br>Machine Check State MSRs."  |
|                  |     | 63:0                                     |                            |                                | Registers R8-15 (and the associated state-save<br>MSRs) exist only in Intel 64 processors. These<br>registers contain valid information only when<br>the processor is operating in 64-bit mode at the<br>time of the error.  |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description   |
|------------------|-----|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec |                                |                    |                            |   |
| 194H             | 404 | MSR_MCG_R12                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check R12<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Registers R8-15 (and the associated state-save MSRs) exist only in Intel 64 processors. These registers contain valid information only when the processor is operating in 64-bit mode at the time of the error. |
| 195H             | 405 | MSR_MCG_R13                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check R13<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Registers R8-15 (and the associated state-save MSRs) exist only in Intel 64 processors. These registers contain valid information only when the processor is operating in 64-bit mode at the time of the error. |
| 196H             | 406 | MSR_MCG_R14                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check R14<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Registers R8-15 (and the associated state-save MSRs) exist only in Intel 64 processors. These registers contain valid information only when the processor is operating in 64-bit mode at the time of the error. |
| 197H             | 407 | MSR_MCG_R15                    | 0, 1, 2, 3, 4, 6   | Unique                     | Machine Check R15<br>See Section 15.3.2.6, "IA32_MCG Extended Machine Check State MSRs."  |
|                  |     | 63:0                           |                    |                            | Registers R8-15 (and the associated state-save MSRs) exist only in Intel 64 processors. These registers contain valid information only when the processor is operating in 64-bit mode at the time of the error. |
| 198H             | 408 | IA32_PERF_STATUS               | 3, 4, 6            | Unique                     | See Table 2-2. See Section 14.1, "Enhanced Intel Speedstep® Technology."  |
| 199H             | 409 | IA32_PERF_CTL                  | 3, 4, 6            | Unique                     | See Table 2-2. See Section 14.1, "Enhanced Intel Speedstep® Technology."  |
| 19AH             | 410 | IA32_CLOCK_MODULATION          | 0, 1, 2, 3, 4, 6   | Unique                     | Thermal Monitor Control (R/W)<br>See Table 2-2.<br>See Section 14.7.3, "Software Controlled Clock Modulation."  |
| 19BH             | 411 | IA32_THERM_INTERRUPT           | 0, 1, 2, 3, 4, 6   | Unique                     | Thermal Interrupt Control (R/W)<br>See Section 14.7.2, "Thermal Monitor," and see Table 2-2.  |
| 19CH             | 412 | IA32_THERM_STATUS              | 0, 1, 2, 3, 4, 6   | Shared                     | Thermal Monitor Status (R/W)<br>See Section 14.7.2, "Thermal Monitor," and see Table 2-2.   |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability   | Shared/<br>Unique <sup>7</sup> | Bit Description  |
|------------------|-----|-----------------------------------|--|--------------------------------|--|
| Hex              | Dec |                                   |  |                                |  |
| 19DH             | 413 | MSR_THERM2_CTL                    |  |                                | Thermal Monitor 2 Control  |
|                  |     |                                   | 3,   | Shared                         | For Family F, Model 3 processors: When read, specifies the value of the target TM2 transition last written. When set, it sets the next target value for TM2 transition.  |
|                  |     |                                   | 4, 6   | Shared                         | For Family F, Model 4 and Model 6 processors: When read, specifies the value of the target TM2 transition last written. Writes may cause #GP exceptions.   |
| 1A0H             | 416 | IA32_MISC_ENABLE                  | 0, 1, 2, 3,<br>4, 6  | Shared                         | Enable Miscellaneous Processor Features (R/W)  |
|                  |     | 0                                 |  |                                | Fast-Strings Enable. See Table 2-2.  |
|                  |     | 1                                 |  |                                | Reserved   |
|                  |     | 2                                 |  |                                | x87 FPU Fopcode Compatibility Mode Enable  |
|                  |     | 3                                 |  |                                | Thermal Monitor 1 Enable<br>See Section 14.7.2, "Thermal Monitor," and see Table 2-2.  |
|                  |     | 4                                 |  |                                | Split-Lock Disable<br>When set, the bit causes an #AC exception to be issued instead of a split-lock cycle. Operating systems that set this bit must align system structures to avoid split-lock scenarios.<br>When the bit is clear (default), normal split-locks are issued to the bus.  |
|                  |     |                                   |  |                                | This debug feature is specific to the Pentium 4 processor.   |
|                  |     | 5                                 |  |                                | Reserved   |
|                  |     | 6                                 |  |                                | Third-Level Cache Disable (R/W)<br>When set, the third-level cache is disabled; when clear (default) the third-level cache is enabled. This flag is reserved for processors that do not have a third-level cache.<br><br>Note that the bit controls only the third-level cache; and only if overall caching is enabled through the CD flag of control register CR0, the page-level cache controls, and/or the MTRRs.<br>See Section 11.5.4, "Disabling and Enabling the L3 Cache." |
|                  |     | 7                                 |  |                                | Performance Monitoring Available (R)<br>See Table 2-2.   |
| 8                |     |                                   | Suppress Lock Enable<br>When set, assertion of LOCK on the bus is suppressed during a Split Lock access. When clear (default), LOCK is not suppressed. |                                |  |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description   |
|------------------|-----|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec |                                |                    |                            |   |
|                  |     | 9                              |                    |                            | Prefetch Queue Disable<br>When set, disables the prefetch queue. When clear (default), enables the prefetch queue.  |
|                  |     | 10                             |                    |                            | FERR# Interrupt Reporting Enable (R/W)<br>When set, interrupt reporting through the FERR# pin is enabled; when clear, this interrupt reporting function is disabled.  |
|                  |     |                                |                    |                            | When this flag is set and the processor is in the stop-clock state (STPCLK# is asserted), asserting the FERR# pin signals to the processor that an interrupt (such as, INIT#, BINIT#, INTR, NMI, SMI#, or RESET#) is pending and that the processor should return to normal operation to handle the interrupt.  |
|                  |     |                                |                    |                            | This flag does not affect the normal operation of the FERR# pin (to indicate an unmasked floating-point error) when the STPCLK# pin is not asserted.  |
|                  |     | 11                             |                    |                            | Branch Trace Storage Unavailable (BTS_UNAVILABLE) (R)<br>See Table 2-2.<br>When set, the processor does not support branch trace storage (BTS); when clear, BTS is supported.   |
|                  |     | 12                             |                    |                            | PEBS_UNAVILABLE: Processor Event Based Sampling Unavailable (R)<br>See Table 2-2.<br>When set, the processor does not support processor event-based sampling (PEBS); when clear, PEBS is supported.   |
|                  |     | 13                             | 3                  |                            | TM2 Enable (R/W)<br>When this bit is set (1) and the thermal sensor indicates that the die temperature is at the pre-determined threshold, the Thermal Monitor 2 mechanism is engaged. TM2 will reduce the bus to core ratio and voltage according to the value last written to MSR_THERM2_CTL bits 15:0.<br>When this bit is clear (0, default), the processor does not change the VID signals or the bus to core ratio when the processor enters a thermal managed state.<br>If the TM2 feature flag (ECX[8]) is not set to 1 after executing CPUID with EAX = 1, then this feature is not supported and BIOS must not alter the contents of this bit location. The processor is operating out of spec if both this bit and the TM1 bit are set to disabled states. |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>7</sup> | Bit Description   |
|------------------|-----|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec |                                |                    |                            |   |
|                  |     | 17:14                          |                    |                            | Reserved  |
|                  |     | 18                             | 3, 4, 6            |                            | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.  |
|                  |     | 19                             |                    |                            | Adjacent Cache Line Prefetch Disable (R/W)<br>When set to 1, the processor fetches the cache line of the 128-byte sector containing currently required data. When set to 0, the processor fetches both cache lines in the sector.   |
|                  |     |                                |                    |                            | Single processor platforms should not set this bit. Server platforms should set or clear this bit based on platform performance observed in validation and testing.<br>BIOS may contain a setup option that controls the setting of this bit.   |
|                  |     | 21:20                          |                    |                            | Reserved  |
|                  |     | 22                             | 3, 4, 6            |                            | Limit CPUID MAXVAL (R/W)<br>See Table 2-2.<br>Setting this can cause unexpected behavior to software that depends on the availability of CPUID leaves greater than 3.   |
|                  |     | 23                             |                    | Shared                     | xTPR Message Disable (R/W)<br>See Table 2-2.  |
|                  |     | 24                             |                    |                            | L1 Data Cache Context Mode (R/W)<br>When set, the L1 data cache is placed in shared mode; when clear (default), the cache is placed in adaptive mode. This bit is only enabled for IA-32 processors that support Intel Hyper-Threading Technology. See Section 11.5.6, "L1 Data Cache Context Mode."<br>When L1 is running in adaptive mode and CR3s are identical, data in L1 is shared across logical processors. Otherwise, L1 is not shared and cache use is competitive.<br>If the Context ID feature flag (ECX[10]) is set to 0 after executing CPUID with EAX = 1, the ability to switch modes is not supported. BIOS must not alter the contents of IA32_MISC_ENABLE[24]. |
|                  |     | 33:25                          |                    |                            | Reserved  |
|                  |     | 34                             |                    | Unique                     | XD Bit Disable (R/W)<br>See Table 2-2.  |
|                  |     | 63:35                          |                    |                            | Reserved  |
| 1A1H             | 417 | MSR_PLATFORM_BRV               | 3, 4, 6            | Shared                     | Platform Feature Requirements (R)   |
|                  |     | 17:0                           |                    |                            | Reserved  |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description  |
|------------------|-----|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec |                                |                    |                            |  |
|                  |     | 18                             |                    |                            | PLATFORM Requirements<br>When set to 1, indicates the processor has specific platform requirements. The details of the platform requirements are listed in the respective data sheets of the processor.  |
|                  |     | 63:19                          |                    |                            | Reserved   |
| 1D7H             | 471 | MSR_LER_FROM_LIP               | 0, 1, 2, 3, 4, 6   | Unique                     | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.<br>See Section 17.13.3, "Last Exception Records."   |
|                  |     | 31:0                           |                    |                            | From Linear IP<br>Linear address of the last branch instruction.   |
|                  |     | 63:32                          |                    |                            | Reserved   |
| 1D7H             | 471 | 63:0                           |                    | Unique                     | From Linear IP<br>Linear address of the last branch instruction (If IA-32e mode is active).  |
| 1D8H             | 472 | MSR_LER_TO_LIP                 | 0, 1, 2, 3, 4, 6   | Unique                     | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.<br>See Section 17.13.3, "Last Exception Records."   |
|                  |     | 31:0                           |                    |                            | From Linear IP<br>Linear address of the target of the last branch instruction.   |
|                  |     | 63:32                          |                    |                            | Reserved   |
| 1D8H             | 472 | 63:0                           |                    | Unique                     | From Linear IP<br>Linear address of the target of the last branch instruction (If IA-32e mode is active).  |
| 1D9H             | 473 | MSR_DEBUGCTLA                  | 0, 1, 2, 3, 4, 6   | Unique                     | Debug Control (R/W)<br>Controls how several debug features are used. Bit definitions are discussed in the referenced section.<br>See Section 17.13.1, "MSR_DEBUGCTLA MSR."   |
| 1DAH             | 474 | MSR_LASTBRANCH_TOS             | 0, 1, 2, 3, 4, 6   | Unique                     | Last Branch Record Stack TOS (R/O)<br>Contains an index (0-3 or 0-15) that points to the top of the last branch record stack (that is, that points the index of the MSR containing the most recent branch record).<br>See Section 17.13.2, "LBR Stack for Processors Based on Intel NetBurst® Microarchitecture"; and addresses 1DBH-1DEH and 680H-68FH. |



Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description   |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|---|
| Hex              | Dec |                                   |                            |                                |   |
| 1DBH             | 475 | MSR_LASTBRANCH_0                  | 0, 1, 2                    | Unique                         | Last Branch Record 0 (R/O)<br>One of four last branch record registers on the last branch record stack. It contains pointers to the source and destination instruction for one of the last four branches, exceptions, or interrupts that the processor took.<br>MSR_LASTBRANCH_0 through MSR_LASTBRANCH_3 at 1DBH-1DEH are available only on family 0FH, models 0H-02H. They have been replaced by the MSRs at 680H-68FH and 6C0H-6CFH. |
|                  |     |                                   |                            |                                | See Section 17.12, "Last Branch, Call Stack, Interrupt, and Exception Recording for Processors based on Skylake Microarchitecture."   |
| 1DCH             | 477 | MSR_LASTBRANCH_1                  | 0, 1, 2                    | Unique                         | Last Branch Record 1<br>See description of the MSR_LASTBRANCH_0 MSR at 1DBH.  |
| 1DDH             | 477 | MSR_LASTBRANCH_2                  | 0, 1, 2                    | Unique                         | Last Branch Record 2<br>See description of the MSR_LASTBRANCH_0 MSR at 1DBH.  |
| 1DEH             | 478 | MSR_LASTBRANCH_3                  | 0, 1, 2                    | Unique                         | Last Branch Record 3<br>See description of the MSR_LASTBRANCH_0 MSR at 1DBH.  |
| 200H             | 512 | IA32_MTRR_PHYSBASE0               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Base MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 201H             | 513 | IA32_MTRR_PHYSMASK0               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 202H             | 514 | IA32_MTRR_PHYSBASE1               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 203H             | 515 | IA32_MTRR_PHYSMASK1               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 204H             | 516 | IA32_MTRR_PHYSBASE2               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 205H             | 517 | IA32_MTRR_PHYSMASK2               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 206H             | 518 | IA32_MTRR_PHYSBASE3               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 207H             | 519 | IA32_MTRR_PHYSMASK3               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 208H             | 520 | IA32_MTRR_PHYSBASE4               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |
| 209H             | 521 | IA32_MTRR_PHYSMASK4               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."  |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>1</sup> | Bit Description  |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|--|
| Hex              | Dec |                                   |                            |                                |  |
| 20AH             | 522 | IA32_MTRR_PHYSBASE5               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 20BH             | 523 | IA32_MTRR_PHYSMASK5               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 20CH             | 524 | IA32_MTRR_PHYSBASE6               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 20DH             | 525 | IA32_MTRR_PHYSMASK6               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 20EH             | 526 | IA32_MTRR_PHYSBASE7               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 20FH             | 527 | IA32_MTRR_PHYSMASK7               | 0, 1, 2, 3,<br>4, 6        | Shared                         | Variable Range Mask MTRR<br>See Section 11.11.2.3, "Variable Range MTRRs."                       |
| 250H             | 592 | IA32_MTRR_FIX64K_00000            | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 258H             | 600 | IA32_MTRR_FIX16K_80000            | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 259H             | 601 | IA32_MTRR_FIX16K_A0000            | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 268H             | 616 | IA32_MTRR_FIX4K_C0000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 269H             | 617 | IA32_MTRR_FIX4K_C8000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26AH             | 618 | IA32_MTRR_FIX4K_D0000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26BH             | 619 | IA32_MTRR_FIX4K_D8000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26CH             | 620 | IA32_MTRR_FIX4K_E0000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26DH             | 621 | IA32_MTRR_FIX4K_E8000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26EH             | 622 | IA32_MTRR_FIX4K_F0000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 26FH             | 623 | IA32_MTRR_FIX4K_F8000             | 0, 1, 2, 3,<br>4, 6        | Shared                         | Fixed Range MTRR<br>See Section 11.11.2.2, "Fixed Range MTRRs."                                  |
| 277H             | 631 | IA32_PAT                          | 0, 1, 2, 3,<br>4, 6        | Unique                         | Page Attribute Table<br>See Section 11.11.2.2, "Fixed Range MTRRs."                              |
| 2FFH             | 767 | IA32_MTRR_DEF_TYPE                | 0, 1, 2, 3,<br>4, 6        | Shared                         | Default Memory Types (R/W)<br>See Table 2-2.<br>See Section 11.11.2.1, "IA32_MTRR_DEF_TYPE MSR." |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description                               |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|---|
| Hex              | Dec |                                   |                            |                                |   |
| 300H             | 768 | MSR_BPU_COUNTER0                  | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 301H             | 769 | MSR_BPU_COUNTER1                  | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 302H             | 770 | MSR_BPU_COUNTER2                  | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 303H             | 771 | MSR_BPU_COUNTER3                  | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 304H             | 772 | MSR_MS_COUNTER0                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 305H             | 773 | MSR_MS_COUNTER1                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 306H             | 774 | MSR_MS_COUNTER2                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 307H             | 775 | MSR_MS_COUNTER3                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 308H             | 776 | MSR_FLAME_COUNTER0                | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 309H             | 777 | MSR_FLAME_COUNTER1                | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30AH             | 778 | MSR_FLAME_COUNTER2                | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30BH             | 779 | MSR_FLAME_COUNTER3                | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30CH             | 780 | MSR_IQ_COUNTER0                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30DH             | 781 | MSR_IQ_COUNTER1                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30EH             | 782 | MSR_IQ_COUNTER2                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 30FH             | 783 | MSR_IQ_COUNTER3                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 310H             | 784 | MSR_IQ_COUNTER4                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 311H             | 785 | MSR_IQ_COUNTER5                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.2, "Performance Counters." |
| 360H             | 864 | MSR_BPU_CCCR0                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs."            |
| 361H             | 865 | MSR_BPU_CCCR1                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs."            |
| 362H             | 866 | MSR_BPU_CCCR2                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs."            |
| 363H             | 867 | MSR_BPU_CCCR3                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs."            |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |     | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>1</sup> | Bit Description                    |
|------------------|-----|-----------------------------------|----------------------------|--------------------------------|------------------------------------|
| Hex              | Dec |                                   |                            |                                |                                    |
| 364H             | 868 | MSR_MS_CCCR0                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 365H             | 869 | MSR_MS_CCCR1                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 366H             | 870 | MSR_MS_CCCR2                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 367H             | 871 | MSR_MS_CCCR3                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 368H             | 872 | MSR_FLAME_CCCR0                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 369H             | 873 | MSR_FLAME_CCCR1                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36AH             | 874 | MSR_FLAME_CCCR2                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36BH             | 875 | MSR_FLAME_CCCR3                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36CH             | 876 | MSR_IQ_CCCR0                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36DH             | 877 | MSR_IQ_CCCR1                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36EH             | 878 | MSR_IQ_CCCR2                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 36FH             | 879 | MSR_IQ_CCCR3                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 370H             | 880 | MSR_IQ_CCCR4                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 371H             | 881 | MSR_IQ_CCCR5                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.3, "CCCR MSRs." |
| 3A0H             | 928 | MSR_BSU_ESCR0                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A1H             | 929 | MSR_BSU_ESCR1                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A2H             | 930 | MSR_FSB_ESCR0                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A3H             | 931 | MSR_FSB_ESCR1                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A4H             | 932 | MSR_FIRM_ESCR0                    | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A5H             | 933 | MSR_FIRM_ESCR1                    | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A6H             | 934 | MSR_FLAME_ESCR0                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |
| 3A7H             | 935 | MSR_FLAME_ESCR1                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 18.6.3.1, "ESCR MSRs." |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |     | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>7</sup> | Bit Description  |
|------------------|-----|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec |                                |                    |                            |  |
| 3A8H             | 936 | MSR_DAC_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3A9H             | 937 | MSR_DAC_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3AAH             | 938 | MSR_MOB_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3ABH             | 939 | MSR_MOB_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3ACH             | 940 | MSR_PMH_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3ADH             | 941 | MSR_PMH_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3AEH             | 942 | MSR_SAA_T_ESCR0                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3AFH             | 943 | MSR_SAA_T_ESCR1                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B0H             | 944 | MSR_U2L_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B1H             | 945 | MSR_U2L_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B2H             | 946 | MSR_BPU_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B3H             | 947 | MSR_BPU_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B4H             | 948 | MSR_IS_ESCR0                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B5H             | 949 | MSR_IS_ESCR1                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B6H             | 950 | MSR_ITLB_ESCR0                 | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B7H             | 951 | MSR_ITLB_ESCR1                 | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B8H             | 952 | MSR_CRU_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3B9H             | 953 | MSR_CRU_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3BAH             | 954 | MSR_IQ_ESCR0                   | 0, 1, 2            | Shared                     | See Section 18.6.3.1, "ESCR MSRs."<br>This MSR is not available on later processors. It is only available on processor family 0FH, models 01H-02H. |
| 3BBH             | 955 | MSR_IQ_ESCR1                   | 0, 1, 2            | Shared                     | See Section 18.6.3.1, "ESCR MSRs."<br>This MSR is not available on later processors. It is only available on processor family 0FH, models 01H-02H. |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |      | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description  |
|------------------|------|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec  |                                |                    |                            |  |
| 3BCH             | 956  | MSR_RAT_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3BDH             | 957  | MSR_RAT_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3BEH             | 958  | MSR_SSU_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3COH             | 960  | MSR_MS_ESCR0                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C1H             | 961  | MSR_MS_ESCR1                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C2H             | 962  | MSR_TBPU_ESCR0                 | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C3H             | 963  | MSR_TBPU_ESCR1                 | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C4H             | 964  | MSR_TC_ESCR0                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C5H             | 965  | MSR_TC_ESCR1                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C8H             | 968  | MSR_IX_ESCR0                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3C9H             | 969  | MSR_IX_ESCR1                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3CAH             | 970  | MSR_ALF_ESCR0                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3CBH             | 971  | MSR_ALF_ESCR1                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3CCH             | 972  | MSR_CRU_ESCR2                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3CDH             | 973  | MSR_CRU_ESCR3                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3E0H             | 992  | MSR_CRU_ESCR4                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3E1H             | 993  | MSR_CRU_ESCR5                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3FOH             | 1008 | MSR_TC_PRECISE_EVENT           | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 18.6.3.1, "ESCR MSRs."   |
| 3F1H             | 1009 | MSR_PEBS_ENABLE                | 0, 1, 2, 3, 4, 6   | Shared                     | Processor Event Based Sampling (PEBS) (R/W)<br>Controls the enabling of processor event sampling and replay tagging. |
|                  |      | 12:0                           |                    |                            | See Table 19-36.   |
|                  |      | 23:13                          |                    |                            | Reserved   |
|                  |      | 24                             |                    |                            | UOP Tag<br>Enables replay tagging when set.  |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |      | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>7</sup> | Bit Description   |
|------------------|------|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec  |                                |                    |                            |   |
|                  |      | 25                             |                    |                            | ENABLE_PEBS_MY_THR (R/W)<br>Enables PEBS for the target logical processor when set; disables PEBS when clear (default). See Section 18.6.4.3, "IA32_PEBS_ENABLE MSR," for an explanation of the target logical processor.<br>This bit is called ENABLE_PEBS in IA-32 processors that do not support Intel Hyper-Threading Technology. |
|                  |      | 26                             |                    |                            | ENABLE_PEBS_OTH_THR (R/W)<br>Enables PEBS for the target logical processor when set; disables PEBS when clear (default). See Section 18.6.4.3, "IA32_PEBS_ENABLE MSR," for an explanation of the target logical processor.<br>This bit is reserved for IA-32 processors that do not support Intel Hyper-Threading Technology.         |
|                  |      | 63:27                          |                    |                            | Reserved  |
| 3F2H             | 1010 | MSR_PEBS_MATRIX_VERT           | 0, 1, 2, 3, 4, 6   | Shared                     | See Table 19-36.  |
| 400H             | 1024 | IA32_MCO_CTL                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 401H             | 1025 | IA32_MCO_STATUS                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 402H             | 1026 | IA32_MCO_ADDR                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.                         |
| 403H             | 1027 | IA32_MCO_MISC                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>The IA32_MCO_MISC MSR is either not implemented or does not contain additional information if the MISC_V flag in the IA32_MCO_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.          |
| 404H             | 1028 | IA32_MC1_CTL                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |      | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description  |
|------------------|------|--------------------------------|--------------------|----------------------------|--|
| Hex              | Dec  |                                |                    |                            |  |
| 406H             | 1030 | IA32_MC1_ADDR                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs." The IA32_MC1_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC1_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.                |
| 407H             | 1031 | IA32_MC1_MISC                  |                    | Shared                     | See Section 15.3.2.4, "IA32_MCI_MISC MSRs." The IA32_MC1_MISC MSR is either not implemented or does not contain additional information if the MISC_V flag in the IA32_MC1_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 408H             | 1032 | IA32_MC2_CTL                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.1, "IA32_MCI_CTL MSRs."   |
| 409H             | 1033 | IA32_MC2_STATUS                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.2, "IA32_MCI_STATUS MSRS."  |
| 40AH             | 1034 | IA32_MC2_ADDR                  |                    |                            | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs." The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.                |
| 40BH             | 1035 | IA32_MC2_MISC                  |                    |                            | See Section 15.3.2.4, "IA32_MCI_MISC MSRs." The IA32_MC2_MISC MSR is either not implemented or does not contain additional information if the MISC_V flag in the IA32_MC2_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40CH             | 1036 | IA32_MC3_CTL                   | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.1, "IA32_MCI_CTL MSRs."   |
| 40DH             | 1037 | IA32_MC3_STATUS                | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.2, "IA32_MCI_STATUS MSRS."  |
| 40EH             | 1038 | IA32_MC3_ADDR                  | 0, 1, 2, 3, 4, 6   | Shared                     | See Section 15.3.2.3, "IA32_MCI_ADDR MSRs." The IA32_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC3_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.                |



Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |      | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description   |
|------------------|------|-----------------------------------|----------------------------|--------------------------------|---|
| Hex              | Dec  |                                   |                            |                                |   |
| 40FH             | 1039 | IA32_MC3_MISC                     | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>The IA32_MC3_MISC MSR is either not implemented or does not contain additional information if the MISCV flag in the IA32_MC3_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 410H             | 1040 | IA32_MC4_CTL                      | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 411H             | 1041 | IA32_MC4_STATUS                   | 0, 1, 2, 3,<br>4, 6        | Shared                         | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."   |
| 412H             | 1042 | IA32_MC4_ADDR                     |                            |                                | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDRIV flag in the IA32_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.               |
| 413H             | 1043 | IA32_MC4_MISC                     |                            |                                | See Section 15.3.2.4, "IA32_MCi_MISC MSRs."<br>The IA32_MC2_MISC MSR is either not implemented or does not contain additional information if the MISCV flag in the IA32_MC4_STATUS register is clear.<br>When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 480H             | 1152 | IA32_VMX_BASIC                    | 3, 4, 6                    | Unique                         | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information."  |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL             | 3, 4, 6                    | Unique                         | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Table 2-2.<br>See Appendix A.3, "VM-Execution Controls."  |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL            | 3, 4, 6                    | Unique                         | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls," and see Table 2-2.   |
| 483H             | 1155 | IA32_VMX_EXIT_CTL                 | 3, 4, 6                    | Unique                         | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Appendix A.4, "VM-Exit Controls," and see Table 2-2.   |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |      | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>1</sup> | Bit Description  |
|------------------|------|-----------------------------------|----------------------------|--------------------------------|--|
| Hex              | Dec  |                                   |                            |                                |  |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL5               | 3, 4, 6                    | Unique                         | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Appendix A.5, "VM-Entry Controls," and see Table 2-2.  |
| 485H             | 1157 | IA32_VMX_MISC                     | 3, 4, 6                    | Unique                         | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Appendix A.6, "Miscellaneous Data," and see Table 2-2.   |
| 486H             | 1158 | IA32_VMX_CR0_FIXED0               | 3, 4, 6                    | Unique                         | Capability Reporting Register of CR0 Bits Fixed to 0 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CR0," and see Table 2-2.  |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1               | 3, 4, 6                    | Unique                         | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CR0," and see Table 2-2.  |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0               | 3, 4, 6                    | Unique                         | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4," and see Table 2-2.  |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1               | 3, 4, 6                    | Unique                         | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4," and see Table 2-2.  |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM                | 3, 4, 6                    | Unique                         | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Appendix A.9, "VMCS Enumeration," and see Table 2-2.  |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTL52          | 3, 4, 6                    | Unique                         | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls," and see Table 2-2.  |
| 600H             | 1536 | IA32_DS_AREA                      | 0, 1, 2, 3, 4, 6           | Unique                         | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."  |
| 680H             | 1664 | MSR_LASTBRANCH_0_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 0 (R/W)<br>One of 16 pairs of last branch record registers on the last branch record stack (680H-68FH). This part of the stack contains pointers to the source instruction for one of the last 16 branches, exceptions, or interrupts taken by the processor. |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |      | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description  |
|------------------|------|-----------------------------------|----------------------------|--------------------------------|--|
| Hex              | Dec  |                                   |                            |                                |  |
|                  |      |                                   |                            |                                | The MSRs at 680H-68FH, 6C0H-6CfH are not available in processor releases before family 0FH, model 03H. These MSRs replace MSRs previously located at 1DBH-1DEH, which performed the same function for early releases.<br><br>See Section 17.12, "Last Branch, Call Stack, Interrupt, and Exception Recording for Processors based on Skylake Microarchitecture." |
| 681H             | 1665 | MSR_LASTBRANCH_1_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 1<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 682H             | 1666 | MSR_LASTBRANCH_2_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 2<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 683H             | 1667 | MSR_LASTBRANCH_3_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 3<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 684H             | 1668 | MSR_LASTBRANCH_4_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 4<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 685H             | 1669 | MSR_LASTBRANCH_5_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 5<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 686H             | 1670 | MSR_LASTBRANCH_6_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 6<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 687H             | 1671 | MSR_LASTBRANCH_7_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 7<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 688H             | 1672 | MSR_LASTBRANCH_8_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 8<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 689H             | 1673 | MSR_LASTBRANCH_9_FROM_IP          | 3, 4, 6                    | Unique                         | Last Branch Record 9<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 68AH             | 1674 | MSR_LASTBRANCH_10_FROM_IP         | 3, 4, 6                    | Unique                         | Last Branch Record 10<br>See description of MSR_LASTBRANCH_0 at 680H.  |
| 68BH             | 1675 | MSR_LASTBRANCH_11_FROM_IP         | 3, 4, 6                    | Unique                         | Last Branch Record 11<br>See description of MSR_LASTBRANCH_0 at 680H.  |
| 68CH             | 1676 | MSR_LASTBRANCH_12_FROM_IP         | 3, 4, 6                    | Unique                         | Last Branch Record 12<br>See description of MSR_LASTBRANCH_0 at 680H.  |

**Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)**

| Register Address |      | Register Name Fields and Flags | Model Availability | Shared/Unique <sup>1</sup> | Bit Description   |
|------------------|------|--------------------------------|--------------------|----------------------------|---|
| Hex              | Dec  |                                |                    |                            |   |
| 68DH             | 1677 | MSR_LASTBRANCH_13_FROM_IP      | 3, 4, 6            | Unique                     | Last Branch Record 13<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 68EH             | 1678 | MSR_LASTBRANCH_14_FROM_IP      | 3, 4, 6            | Unique                     | Last Branch Record 14<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 68FH             | 1679 | MSR_LASTBRANCH_15_FROM_IP      | 3, 4, 6            | Unique                     | Last Branch Record 15<br>See description of MSR_LASTBRANCH_0 at 680H.   |
| 6C0H             | 1728 | MSR_LASTBRANCH_0_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 0 (R/W)<br>One of 16 pairs of last branch record registers on the last branch record stack (6C0H-6CFH). This part of the stack contains pointers to the destination instruction for one of the last 16 branches, exceptions, or interrupts that the processor took.<br>See Section 17.12, "Last Branch, Call Stack, Interrupt, and Exception Recording for Processors based on Skylake Microarchitecture." |
| 6C1H             | 1729 | MSR_LASTBRANCH_1_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 1<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C2H             | 1730 | MSR_LASTBRANCH_2_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 2<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C3H             | 1731 | MSR_LASTBRANCH_3_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 3<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C4H             | 1732 | MSR_LASTBRANCH_4_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 4<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C5H             | 1733 | MSR_LASTBRANCH_5_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 5<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C6H             | 1734 | MSR_LASTBRANCH_6_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 6<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C7H             | 1735 | MSR_LASTBRANCH_7_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 7<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |
| 6C8H             | 1736 | MSR_LASTBRANCH_8_TO_IP         | 3, 4, 6            | Unique                     | Last Branch Record 8<br>See description of MSR_LASTBRANCH_0 at 6C0H.  |

Table 2-45. MSRs in the Pentium® 4 and Intel® Xeon® Processors (Contd.)

| Register Address |      | Register Name<br>Fields and Flags | Model<br>Avail-<br>ability | Shared/<br>Unique <sup>7</sup> | Bit Description   |
|------------------|------|-----------------------------------|----------------------------|--------------------------------|---|
| Hex              | Dec  |                                   |                            |                                |   |
| 6C9H             | 1737 | MSR_LASTBRANCH_9_TO_IP            | 3, 4, 6                    | Unique                         | Last Branch Record 9<br>See description of MSR_LASTBRANCH_0 at 6COH.  |
| 6CAH             | 1738 | MSR_LASTBRANCH_10_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 10<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| 6CBH             | 1739 | MSR_LASTBRANCH_11_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 11<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| 6CCH             | 1740 | MSR_LASTBRANCH_12_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 12<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| 6CDH             | 1741 | MSR_LASTBRANCH_13_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 13<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| 6CEH             | 1742 | MSR_LASTBRANCH_14_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 14<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| 6CFH             | 1743 | MSR_LASTBRANCH_15_TO_IP           | 3, 4, 6                    | Unique                         | Last Branch Record 15<br>See description of MSR_LASTBRANCH_0 at 6COH. |
| C000_<br>0080H   |      | IA32_EFER                         | 3, 4, 6                    | Unique                         | Extended Feature Enables<br>See Table 2-2.                            |
| C000_<br>0081H   |      | IA32_STAR                         | 3, 4, 6                    | Unique                         | System Call Target Address (R/W)<br>See Table 2-2.                    |
| C000_<br>0082H   |      | IA32_LSTAR                        | 3, 4, 6                    | Unique                         | IA-32e Mode System Call Target Address (R/W)<br>See Table 2-2.        |
| C000_<br>0084H   |      | IA32_FMASK                        | 3, 4, 6                    | Unique                         | System Call Flag Mask (R/W)<br>See Table 2-2.                         |
| C000_<br>0100H   |      | IA32_FS_BASE                      | 3, 4, 6                    | Unique                         | Map of BASE Address of FS (R/W)<br>See Table 2-2.                     |
| C000_<br>0101H   |      | IA32_GS_BASE                      | 3, 4, 6                    | Unique                         | Map of BASE Address of GS (R/W)<br>See Table 2-2.                     |
| C000_<br>0102H   |      | IA32_KERNEL_GS_BASE               | 3, 4, 6                    | Unique                         | Swap Target of BASE Address of GS (R/W)<br>See Table 2-2.             |

**NOTES**

1. For HT-enabled processors, there may be more than one logical processors per physical unit. If an MSR is Shared, this means that one MSR is shared between logical processors. If an MSR is unique, this means that each logical processor has its own MSR.

## 2.18.1 MSRs Unique to Intel® Xeon® Processor MP with L3 Cache

The MSRs listed in Table 2-46 apply to Intel® Xeon® Processor MP with up to 8MB level three cache. These processors can be detected by enumerating the deterministic cache parameter leaf of CPUID instruction (with EAX = 4 as input) to detect the presence of the third level cache, and with CPUID reporting family encoding 0FH, model encoding 3 or 4 (see CPUID instruction for more details).

**Table 2-46. MSRs Unique to 64-bit Intel® Xeon® Processor MP with Up to an 8 MB L3 Cache**

| Register Address | Register Name Fields and Flags | Model Availability | Shared/Unique | Bit Description   |
|------------------|--------------------------------|--------------------|---------------|---|
| 107CCH           | MSR_IFSB_BUSQ0                 | 3, 4               | Shared        | IFSB BUSQ Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache." |
| 107CDH           | MSR_IFSB_BUSQ1                 | 3, 4               | Shared        | IFSB BUSQ Event Control and Counter Register (R/W)  |
| 107CEH           | MSR_IFSB_SNPQ0                 | 3, 4               | Shared        | IFSB SNPQ Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache." |
| 107CFH           | MSR_IFSB_SNPQ1                 | 3, 4               | Shared        | IFSB SNPQ Event Control and Counter Register (R/W)  |
| 107D0H           | MSR_EFSB_DRDY0                 | 3, 4               | Shared        | EFSB DRDY Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache." |
| 107D1H           | MSR_EFSB_DRDY1                 | 3, 4               | Shared        | EFSB DRDY Event Control and Counter Register (R/W)  |
| 107D2H           | MSR_IFSB_CTL6                  | 3, 4               | Shared        | IFSB Latency Event Control Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache."          |
| 107D3H           | MSR_IFSB_CNTR7                 | 3, 4               | Shared        | IFSB Latency Event Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache."          |

The MSRs listed in Table 2-47 apply to Intel® Xeon® Processor 7100 series. These processors can be detected by enumerating the deterministic cache parameter leaf of CPUID instruction (with EAX = 4 as input) to detect the presence of the third level cache, and with CPUID reporting family encoding 0FH, model encoding 6 (See CPUID instruction for more details.). The performance monitoring MSRs listed in Table 2-47 are shared between logical processors in the same core, but are replicated for each core.

**Table 2-47. MSRs Unique to Intel® Xeon® Processor 7100 Series**

| Register Address | Register Name Fields and Flags | Model Availability | Shared/Unique | Bit Description   |
|------------------|--------------------------------|--------------------|---------------|---|
| 107CCH           | MSR_EMON_L3_CTR_CTL0           | 6                  | Shared        | GBUSQ Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache." |

**Table 2-47. MSRs Unique to Intel® Xeon® Processor 7100 Series (Contd.)**

| Register Address |  | Register Name<br>Fields and Flags | Model Avail-<br>ability | Shared/<br>Unique | Bit Description   |
|------------------|--|-----------------------------------|-------------------------|-------------------|---|
| 107CDH           |  | MSR_EMON_L3_CTR_CTL1              | 6                       | Shared            | GBUSQ Event Control and Counter Register (R/W)  |
| 107CEH           |  | MSR_EMON_L3_CTR_CTL2              | 6                       | Shared            | GSNPQ Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache." |
| 107CFH           |  | MSR_EMON_L3_CTR_CTL3              | 6                       | Shared            | GSNPQ Event Control and Counter Register (R/W)  |
| 107D0H           |  | MSR_EMON_L3_CTR_CTL4              | 6                       | Shared            | FSB Event Control and Counter Register (R/W)<br>See Section 18.6.6, "Performance Monitoring on 64-bit Intel Xeon Processor MP with Up to 8-MByte L3 Cache."   |
| 107D1H           |  | MSR_EMON_L3_CTR_CTL5              | 6                       | Shared            | FSB Event Control and Counter Register (R/W)  |
| 107D2H           |  | MSR_EMON_L3_CTR_CTL6              | 6                       | Shared            | FSB Event Control and Counter Register (R/W)  |
| 107D3H           |  | MSR_EMON_L3_CTR_CTL7              | 6                       | Shared            | FSB Event Control and Counter Register (R/W)  |

## 2.19 MSRS IN INTEL® CORE™ SOLO AND INTEL® CORE™ DUO PROCESSORS

Model-specific registers (MSRs) for Intel Core Solo, Intel Core Duo processors, and Dual-core Intel Xeon processor LV are listed in Table 2-48. The column "Shared/Unique" applies to Intel Core Duo processor. "Unique" means each processor core has a separate MSR, or a bit field in an MSR governs only a core independently. "Shared" means the MSR or the bit field in an MSR address governs the operation of both processor cores.

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |     | Register Name            | Shared/<br>Unique | Bit Description   |
|------------------|-----|--------------------------|-------------------|---|
| Hex              | Dec |                          |                   |   |
| 0H               | 0   | P5_MC_ADDR               | Unique            | See Section 2.22, "MSRs in Pentium Processors," and see Table 2-2.  |
| 1H               | 1   | P5_MC_TYPE               | Unique            | See Section 2.22, "MSRs in Pentium Processors," and see Table 2-2.  |
| 6H               | 6   | IA32_MONITOR_FILTER_SIZE | Unique            | See Section 8.10.5, "Monitor/Mwait Address Range Determination," and see Table 2-2.   |
| 10H              | 16  | IA32_TIME_STAMP_COUNTER  | Unique            | See Section 17.17, "Time-Stamp Counter," and see Table 2-2.   |
| 17H              | 23  | IA32_PLATFORM_ID         | Shared            | Platform ID (R)<br>See Table 2-2.<br>The operating system can use this MSR to determine "slot" information for the processor and the proper microcode update to load. |
| 1BH              | 27  | IA32_APIC_BASE           | Unique            | See Section 10.4.4, "Local APIC Status and Location," and see Table 2-2.  |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |                             | Register Name      | Shared/<br>Unique | Bit Description  |
|------------------|-----------------------------|--------------------|-------------------|--|
| Hex              | Dec                         |                    |                   |  |
| 2AH              | 42                          | MSR_EBL_CR_POWERON | Shared            | Processor Hard Power-On Configuration (R/W)<br>Enables and disables processor features; (R) indicates current processor configuration. |
|                  |                             | 0                  |                   | Reserved   |
|                  |                             | 1                  |                   | Data Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                               |
|                  |                             | 2                  |                   | Response Error Checking Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                           |
|                  |                             | 3                  |                   | MCERR# Drive Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                                      |
|                  |                             | 4                  |                   | Address Parity Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                                    |
|                  |                             | 6: 5               |                   | Reserved   |
|                  |                             | 7                  |                   | BINIT# Driver Enable (R/W)<br>1 = Enabled; 0 = Disabled<br>Note: Not all processor implements R/W.                                     |
|                  |                             | 8                  |                   | Output Tri-state Enabled (R/O)<br>1 = Enabled; 0 = Disabled  |
|                  |                             | 9                  |                   | Execute BIST (R/O)<br>1 = Enabled; 0 = Disabled  |
|                  |                             | 10                 |                   | MCERR# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled  |
|                  |                             | 11                 |                   | Reserved   |
|                  |                             | 12                 |                   | BINIT# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled  |
|                  |                             | 13                 |                   | Reserved   |
|                  |                             | 14                 |                   | 1 MByte Power on Reset Vector (R/O)<br>1 = 1 MByte; 0 = 4 GBytes   |
|                  |                             | 15                 |                   | Reserved   |
|                  |                             | 17:16              |                   | APIC Cluster ID (R/O)  |
|                  |                             | 18                 |                   | System Bus Frequency (R/O)<br>0 = 100 MHz<br>1 = Reserved  |
|                  |                             | 19                 |                   | Reserved   |
|                  |                             | 21: 20             |                   | Symmetric Arbitration ID (R/O)   |
| 26:22            | Clock Frequency Ratio (R/O) |                    |                   |  |



Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV

| Register Address |     | Register Name        | Shared/<br>Unique | Bit Description   |
|------------------|-----|----------------------|-------------------|---|
| Hex              | Dec |                      |                   |   |
| 3AH              | 58  | IA32_FEATURE_CONTROL | Unique            | Control Features in IA-32 Processor (R/W)<br>See Table 2-2.   |
| 40H              | 64  | MSR_LASTBRANCH_0     | Unique            | Last Branch Record 0 (R/W)<br>One of 8 last branch record registers on the last branch record stack: bits 31-0 hold the 'from' address and bits 63-32 hold the 'to' address. See also: <ul style="list-style-type: none"> <li>Last Branch Record Stack TOS at 1C9H</li> <li>Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)."</li> </ul> |
| 41H              | 65  | MSR_LASTBRANCH_1     | Unique            | Last Branch Record 1 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 42H              | 66  | MSR_LASTBRANCH_2     | Unique            | Last Branch Record 2 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 43H              | 67  | MSR_LASTBRANCH_3     | Unique            | Last Branch Record 3 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 44H              | 68  | MSR_LASTBRANCH_4     | Unique            | Last Branch Record 4 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 45H              | 69  | MSR_LASTBRANCH_5     | Unique            | Last Branch Record 5 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 46H              | 70  | MSR_LASTBRANCH_6     | Unique            | Last Branch Record 6 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 47H              | 71  | MSR_LASTBRANCH_7     | Unique            | Last Branch Record 7 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 79H              | 121 | IA32_BIOS_UPDT_TRIG  | Unique            | BIOS Update Trigger Register (W)<br>See Table 2-2.  |
| 8BH              | 139 | IA32_BIOS_SIGN_ID    | Unique            | BIOS Update Signature ID (RO)<br>See Table 2-2.   |
| C1H              | 193 | IA32_PMC0            | Unique            | Performance Counter Register<br>See Table 2-2.  |
| C2H              | 194 | IA32_PMC1            | Unique            | Performance Counter Register<br>See Table 2-2.  |
| CDH              | 205 | MSR_FSB_FREQ         | Shared            | Scaleable Bus Speed (RO)<br>This field indicates the scaleable bus clock speed:   |
|                  |     | 2:0                  |                   | <ul style="list-style-type: none"> <li>101B: 100 MHz (FSB 400)</li> <li>001B: 133 MHz (FSB 533)</li> <li>011B: 167 MHz (FSB 667)</li> </ul> 133.33 MHz should be utilized if performing calculation with System Bus Speed when encoding is 101B.<br>166.67 MHz should be utilized if performing calculation with System Bus Speed when encoding is 001B.                        |
|                  |     | 63:3                 |                   | Reserved  |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |     | Register Name     | Shared/<br>Unique | Bit Description   |
|------------------|-----|-------------------|-------------------|---|
| Hex              | Dec |                   |                   |   |
| E7H              | 231 | IA32_MPERF        | Unique            | Maximum Performance Frequency Clock Count (RW)<br>See Table 2-2.  |
| E8H              | 232 | IA32_APERF        | Unique            | Actual Performance Frequency Clock Count (RW)<br>See Table 2-2.   |
| FEH              | 254 | IA32_MTRRCAP      | Unique            | See Table 2-2.  |
| 11EH             | 281 | MSR_BBL_CR_CTL3   | Shared            | Control Register 3<br>Used to configure the L2 Cache.   |
|                  |     | 0                 |                   | L2 Hardware Enabled (RO)<br>1 = If the L2 is hardware-enabled<br>0 = Indicates if the L2 is hardware-disabled   |
|                  |     | 7:1               |                   | Reserved  |
|                  |     | 8                 |                   | L2 Enabled (R/W)<br>1 = L2 cache has been initialized<br>0 = Disabled (default)<br>Until this bit is set the processor will not respond to the WBINVD instruction or the assertion of the FLUSH# input.   |
|                  |     | 22:9              |                   | Reserved  |
|                  |     | 23                |                   | L2 Not Present (RO)<br>0 = L2 Present<br>1 = L2 Not Present   |
|                  |     | 63:24             |                   | Reserved  |
| 174H             | 372 | IA32_SYSENTER_CS  | Unique            | See Table 2-2.  |
| 175H             | 373 | IA32_SYSENTER_ESP | Unique            | See Table 2-2.  |
| 176H             | 374 | IA32_SYSENTER_EIP | Unique            | See Table 2-2.  |
| 179H             | 377 | IA32_MCG_CAP      | Unique            | See Table 2-2.  |
| 17AH             | 378 | IA32_MCG_STATUS   | Unique            | Global Machine Check Status   |
|                  |     | 0                 |                   | RIPV<br>When set, this bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If this bit is cleared, the program cannot be reliably restarted. |
|                  |     | 1                 |                   | EIPV<br>When set, this bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |

Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV

| Register Address |     | Register Name         | Shared/<br>Unique | Bit Description   |
|------------------|-----|-----------------------|-------------------|---|
| Hex              | Dec |                       |                   |   |
|                  |     | 2                     |                   | MCIP<br>When set, this bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception.   |
|                  |     | 63:3                  |                   | Reserved  |
| 186H             | 390 | IA32_PERFEVTSELO      | Unique            | See Table 2-2.  |
| 187H             | 391 | IA32_PERFEVTSEL1      | Unique            | See Table 2-2.  |
| 198H             | 408 | IA32_PERF_STATUS      | Shared            | See Table 2-2.  |
| 199H             | 409 | IA32_PERF_CTL         | Unique            | See Table 2-2.  |
| 19AH             | 410 | IA32_CLOCK_MODULATION | Unique            | Clock Modulation (R/W)<br>See Table 2-2.  |
| 19BH             | 411 | IA32_THERM_INTERRUPT  | Unique            | Thermal Interrupt Control (R/W)<br>See Table 2-2.<br>See Section 14.7.2, "Thermal Monitor."   |
| 19CH             | 412 | IA32_THERM_STATUS     | Unique            | Thermal Monitor Status (R/W)<br>See Table 2-2.<br>See Section 14.7.2, "Thermal Monitor".  |
| 19DH             | 413 | MSR_THERM2_CTL        | Unique            | Thermal Monitor 2 Control   |
|                  |     | 15:0                  |                   | Reserved  |
|                  |     | 16                    |                   | TM_SELECT (R/W)<br>Mode of automatic thermal monitor:<br>0 = Thermal Monitor 1 (thermally-initiated on-die modulation of the stop-clock duty cycle)<br>1 = Thermal Monitor 2 (thermally-initiated frequency transitions)<br>If bit 3 of the IA32_MISC_ENABLE register is cleared, TM_SELECT has no effect. Neither TM1 nor TM2 will be enabled. |
|                  |     | 63:16                 |                   | Reserved  |
| 1A0H             | 416 | IA32_MISC_ENABLE      |                   | Enable Miscellaneous Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |     | 2:0                   |                   | Reserved  |
|                  |     | 3                     | Unique            | Automatic Thermal Control Circuit Enable (R/W)<br>See Table 2-2.  |
|                  |     | 6:4                   |                   | Reserved  |
|                  |     | 7                     | Shared            | Performance Monitoring Available (R)<br>See Table 2-2.  |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |     | Register Name      | Shared/<br>Unique | Bit Description   |
|------------------|-----|--------------------|-------------------|---|
| Hex              | Dec |                    |                   |   |
|                  |     | 9:8                |                   | Reserved  |
|                  |     | 10                 | Shared            | FERR# Multiplexing Enable (R/W)<br>1 = FERR# asserted by the processor to indicate a pending break event within the processor<br>0 = Indicates compatible FERR# signaling behavior<br>This bit must be set to 1 to support XAPIC interrupt model usage.   |
|                  |     | 11                 | Shared            | Branch Trace Storage Unavailable (RO)<br>See Table 2-2.   |
|                  |     | 12                 |                   | Reserved  |
|                  |     | 13                 | Shared            | TM2 Enable (R/W)<br>When this bit is set (1) and the thermal sensor indicates that the die temperature is at the pre-determined threshold, the Thermal Monitor 2 mechanism is engaged. TM2 will reduce the bus to core ratio and voltage according to the value last written to MSR_THERM2_CTL bits 15:0.<br>When this bit is clear (0, default), the processor does not change the VID signals or the bus to core ratio when the processor enters a thermal managed state.<br>If the TM2 feature flag (ECX[8]) is not set to 1 after executing CPUID with EAX = 1, then this feature is not supported and BIOS must not alter the contents of this bit location. The processor is operating out of spec if both this bit and the TM1 bit are set to disabled states. |
|                  |     | 15:14              |                   | Reserved  |
|                  |     | 16                 | Shared            | Enhanced Intel SpeedStep Technology Enable (R/W)<br>1 = Enhanced Intel SpeedStep Technology enabled   |
|                  |     | 18                 | Shared            | ENABLE MONITOR FSM (R/W)<br>See Table 2-2.  |
|                  |     | 19                 |                   | Reserved  |
|                  |     | 22                 | Shared            | Limit CPUID Maxval (R/W)<br>See Table 2-2.<br>Setting this bit may cause behavior in software that depends on the availability of CPUID leaves greater than 2.  |
|                  |     | 33:23              |                   | Reserved  |
|                  |     | 34                 | Shared            | XD Bit Disable (R/W)<br>See Table 2-2.  |
|                  |     | 63:35              |                   | Reserved  |
| 1C9H             | 457 | MSR_LASTBRANCH_TOS | Unique            | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-3) that points to the MSR containing the most recent branch record.<br>See MSR_LASTBRANCH_0_FROM_IP (at 40H).   |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |     | Register Name    | Shared/<br>Unique | Bit Description  |
|------------------|-----|------------------|-------------------|--|
| Hex              | Dec |                  |                   |  |
| 1D9H             | 473 | IA32_DEBUGCTL    | Unique            | Debug Control (R/W)<br>Controls how several debug features are used. Bit definitions are discussed in Table 2-2.   |
| 1DDH             | 477 | MSR_LER_FROM_LIP | Unique            | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.                       |
| 1DEH             | 478 | MSR_LER_TO_LIP   | Unique            | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled. |
| 200H             | 512 | MTRRphysBase0    | Unique            | Memory Type Range Registers  |
| 201H             | 513 | MTRRphysMask0    | Unique            | Memory Type Range Registers  |
| 202H             | 514 | MTRRphysBase1    | Unique            | Memory Type Range Registers  |
| 203H             | 515 | MTRRphysMask1    | Unique            | Memory Type Range Registers  |
| 204H             | 516 | MTRRphysBase2    | Unique            | Memory Type Range Registers  |
| 205H             | 517 | MTRRphysMask2    | Unique            | Memory Type Range Registers  |
| 206H             | 518 | MTRRphysBase3    | Unique            | Memory Type Range Registers  |
| 207H             | 519 | MTRRphysMask3    | Unique            | Memory Type Range Registers  |
| 208H             | 520 | MTRRphysBase4    | Unique            | Memory Type Range Registers  |
| 209H             | 521 | MTRRphysMask4    | Unique            | Memory Type Range Registers  |
| 20AH             | 522 | MTRRphysBase5    | Unique            | Memory Type Range Registers  |
| 20BH             | 523 | MTRRphysMask5    | Unique            | Memory Type Range Registers  |
| 20CH             | 524 | MTRRphysBase6    | Unique            | Memory Type Range Registers  |
| 20DH             | 525 | MTRRphysMask6    | Unique            | Memory Type Range Registers  |
| 20EH             | 526 | MTRRphysBase7    | Unique            | Memory Type Range Registers  |
| 20FH             | 527 | MTRRphysMask7    | Unique            | Memory Type Range Registers  |
| 250H             | 592 | MTRRfix64K_00000 | Unique            | Memory Type Range Registers  |
| 258H             | 600 | MTRRfix16K_80000 | Unique            | Memory Type Range Registers  |
| 259H             | 601 | MTRRfix16K_A0000 | Unique            | Memory Type Range Registers  |
| 268H             | 616 | MTRRfix4K_C0000  | Unique            | Memory Type Range Registers  |
| 269H             | 617 | MTRRfix4K_C8000  | Unique            | Memory Type Range Registers  |
| 26AH             | 618 | MTRRfix4K_D0000  | Unique            | Memory Type Range Registers  |
| 26BH             | 619 | MTRRfix4K_D8000  | Unique            | Memory Type Range Registers  |
| 26CH             | 620 | MTRRfix4K_E0000  | Unique            | Memory Type Range Registers  |
| 26DH             | 621 | MTRRfix4K_E8000  | Unique            | Memory Type Range Registers  |
| 26EH             | 622 | MTRRfix4K_F0000  | Unique            | Memory Type Range Registers  |
| 26FH             | 623 | MTRRfix4K_F8000  | Unique            | Memory Type Range Registers  |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |      | Register Name      | Shared/<br>Unique | Bit Description  |
|------------------|------|--------------------|-------------------|--|
| Hex              | Dec  |                    |                   |  |
| 2FFH             | 767  | IA32_MTRR_DEF_TYPE | Unique            | Default Memory Types (R/W)<br>See Table 2-2.<br>See Section 11.11.2.1, "IA32_MTRR_DEF_TYPE MSR."   |
| 400H             | 1024 | IA32_MCO_CTL       | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 401H             | 1025 | IA32_MCO_STATUS    | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 402H             | 1026 | IA32_MCO_ADDR      | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 404H             | 1028 | IA32_MC1_CTL       | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 405H             | 1029 | IA32_MC1_STATUS    | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 406H             | 1030 | IA32_MC1_ADDR      | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC1_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC1_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 408H             | 1032 | IA32_MC2_CTL       | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 409H             | 1033 | IA32_MC2_STATUS    | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 40AH             | 1034 | IA32_MC2_ADDR      | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40CH             | 1036 | MSR_MC4_CTL        | Unique            | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 40DH             | 1037 | MSR_MC4_STATUS     | Unique            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 40EH             | 1038 | MSR_MC4_ADDR       | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 410H             | 1040 | IA32_MC3_CTL       |                   | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 411H             | 1041 | IA32_MC3_STATUS    |                   | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 412H             | 1042 | MSR_MC3_ADDR       | Unique            | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |

Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV

| Register Address |      | Register Name          | Shared/<br>Unique | Bit Description  |
|------------------|------|------------------------|-------------------|--|
| Hex              | Dec  |                        |                   |  |
| 413H             | 1043 | MSR_MC3_MISC           | Unique            | Machine Check Error Reporting Register - contains additional information describing the machine-check error if the MISCV flag in the IA32_MCI_STATUS register is set.  |
| 414H             | 1044 | MSR_MC5_CTL            | Unique            | Machine Check Error Reporting Register - controls signaling of #MC for errors produced by a particular hardware unit (or group of hardware units).   |
| 415H             | 1045 | MSR_MC5_STATUS         | Unique            | Machine Check Error Reporting Register - contains information related to a machine-check error if its VAL (valid) flag is set. Software is responsible for clearing IA32_MCI_STATUS MSRs by explicitly writing 0s to them; writing 1s to them causes a general-protection exception. |
| 416H             | 1046 | MSR_MC5_ADDR           | Unique            | Machine Check Error Reporting Register - contains the address of the code or data memory location that produced the machine-check error if the ADDRV flag in the IA32_MCI_STATUS register is set.  |
| 417H             | 1047 | MSR_MC5_MISC           | Unique            | Machine Check Error Reporting Register - contains additional information describing the machine-check error if the MISCV flag in the IA32_MCI_STATUS register is set.  |
| 480H             | 1152 | IA32_VMX_BASIC         | Unique            | Reporting Register of Basic VMX Capabilities (R/O)<br>See Table 2-2.<br>See Appendix A.1, "Basic VMX Information".<br>(If CPUID.01H:ECX.[bit 9])   |
| 481H             | 1153 | IA32_VMX_PINBASED_CTL  | Unique            | Capability Reporting Register of Pin-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls".<br>(If CPUID.01H:ECX.[bit 9])   |
| 482H             | 1154 | IA32_VMX_PROCBASED_CTL | Unique            | Capability Reporting Register of Primary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls".<br>(If CPUID.01H:ECX.[bit 9])   |
| 483H             | 1155 | IA32_VMX_EXIT_CTL      | Unique            | Capability Reporting Register of VM-Exit Controls (R/O)<br>See Appendix A.4, "VM-Exit Controls".<br>(If CPUID.01H:ECX.[bit 9])   |
| 484H             | 1156 | IA32_VMX_ENTRY_CTL     | Unique            | Capability Reporting Register of VM-Entry Controls (R/O)<br>See Appendix A.5, "VM-Entry Controls".<br>(If CPUID.01H:ECX.[bit 9])   |
| 485H             | 1157 | IA32_VMX_MISC          | Unique            | Reporting Register of Miscellaneous VMX Capabilities (R/O)<br>See Appendix A.6, "Miscellaneous Data".<br>(If CPUID.01H:ECX.[bit 9])  |
| 486H             | 1158 | IA32_VMX_CRO_FIXED0    | Unique            | Capability Reporting Register of CRO Bits Fixed to 0 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CRO".<br>(If CPUID.01H:ECX.[bit 9])   |

**Table 2-48. MSRs in Intel® Core™ Solo, Intel® Core™ Duo Processors, and Dual-Core Intel® Xeon® Processor LV**

| Register Address |      | Register Name             | Shared/<br>Unique | Bit Description   |
|------------------|------|---------------------------|-------------------|---|
| Hex              | Dec  |                           |                   |   |
| 487H             | 1159 | IA32_VMX_CR0_FIXED1       | Unique            | Capability Reporting Register of CR0 Bits Fixed to 1 (R/O)<br>See Appendix A.7, "VMX-Fixed Bits in CR0".<br>(If CPUID.01H:ECX.[bit 9])  |
| 488H             | 1160 | IA32_VMX_CR4_FIXED0       | Unique            | Capability Reporting Register of CR4 Bits Fixed to 0 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4".<br>(If CPUID.01H:ECX.[bit 9])  |
| 489H             | 1161 | IA32_VMX_CR4_FIXED1       | Unique            | Capability Reporting Register of CR4 Bits Fixed to 1 (R/O)<br>See Appendix A.8, "VMX-Fixed Bits in CR4".<br>(If CPUID.01H:ECX.[bit 9])  |
| 48AH             | 1162 | IA32_VMX_VMCS_ENUM        | Unique            | Capability Reporting Register of VMCS Field Enumeration (R/O)<br>See Appendix A.9, "VMCS Enumeration".<br>(If CPUID.01H:ECX.[bit 9])  |
| 48BH             | 1163 | IA32_VMX_PROCBASED_CTLSS2 | Unique            | Capability Reporting Register of Secondary Processor-Based VM-Execution Controls (R/O)<br>See Appendix A.3, "VM-Execution Controls".<br>(If CPUID.01H:ECX.[bit 9] and IA32_VMX_PROCBASED_CTLSS[bit 63]) |
| 600H             | 1536 | IA32_DS_AREA              | Unique            | DS Save Area (R/W)<br>See Table 2-2.<br>See Section 18.6.3.4, "Debug Store (DS) Mechanism."   |
|                  |      | 31:0                      |                   | DS Buffer Management Area<br>Linear address of the first byte of the DS buffer management area.   |
|                  |      | 63:32                     |                   | Reserved  |
| C000_0080H       |      | IA32_EFER                 | Unique            | See Table 2-2.  |
|                  |      | 10:0                      |                   | Reserved  |
|                  |      | 11                        |                   | Execute Disable Bit Enable  |
|                  |      | 63:12                     |                   | Reserved  |

## 2.20 MSRS IN THE PENTIUM M PROCESSOR

Model-specific registers (MSRs) for the Pentium M processor are similar to those described in Section 2.21 for P6 family processors. The following table describes new MSRs and MSRs whose behavior has changed on the Pentium M processor.

**Table 2-49. MSRs in Pentium M Processors**

| Register Address |     | Register Name / Bit Fields | Bit Description                                 |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
| 0H               | 0   | P5_MC_ADDR                 | See Section 2.22, "MSRs in Pentium Processors." |



Table 2-49. MSRs in Pentium M Processors (Contd.)

| Register Address |          | Register Name / Bit Fields | Bit Description   |
|------------------|----------|----------------------------|---|
| Hex              | Dec      |                            |   |
| 1H               | 1        | P5_MC_TYPE                 | See Section 2.22, "MSRs in Pentium Processors."   |
| 10H              | 16       | IA32_TIME_STAMP_COUNTER    | See Section 17.17, "Time-Stamp Counter," and see Table 2-2.   |
| 17H              | 23       | IA32_PLATFORM_ID           | Platform ID (R)<br>See Table 2-2.<br>The operating system can use this MSR to determine "slot" information for the processor and the proper microcode update to load. |
| 2AH              | 42       | MSR_EBL_CR_POWERON         | Processor Hard Power-On Configuration<br>(R/W) Enables and disables processor features.<br>(R) Indicates current processor configuration.                             |
|                  |          | 0                          | Reserved  |
|                  |          | 1                          | Data Error Checking Enable (R)<br>0 = Disabled<br>Always 0 on the Pentium M processor.  |
|                  |          | 2                          | Response Error Checking Enable (R)<br>0 = Disabled<br>Always 0 on the Pentium M processor.  |
|                  |          | 3                          | MCERR# Drive Enable (R)<br>0 = Disabled<br>Always 0 on the Pentium M processor.   |
|                  |          | 4                          | Address Parity Enable (R)<br>0 = Disabled<br>Always 0 on the Pentium M processor.   |
|                  |          | 6:5                        | Reserved  |
|                  |          | 7                          | BINIT# Driver Enable (R)<br>1 = Enabled; 0 = Disabled<br>Always 0 on the Pentium M processor.   |
|                  |          | 8                          | Output Tri-state Enabled (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |          | 9                          | Execute BIST (R/O)<br>1 = Enabled; 0 = Disabled   |
|                  |          | 10                         | MCERR# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled<br>Always 0 on the Pentium M processor.   |
|                  |          | 11                         | Reserved  |
|                  |          | 12                         | BINIT# Observation Enabled (R/O)<br>1 = Enabled; 0 = Disabled<br>Always 0 on the Pentium M processor.   |
| 13               | Reserved |                            |   |

**Table 2-49. MSRs in Pentium M Processors (Contd.)**

| Register Address |     | Register Name / Bit Fields | Bit Description   |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
|                  |     | 14                         | 1 MByte Power on Reset Vector (R/O)<br>1 = 1 MByte; 0 = 4 GBytes<br>Always 0 on the Pentium M processor.  |
|                  |     | 15                         | Reserved  |
|                  |     | 17:16                      | APIC Cluster ID (R/O)<br>Always 00B on the Pentium M processor.   |
|                  |     | 18                         | System Bus Frequency (R/O)<br>0 = 100 MHz<br>1 = Reserved<br>Always 0 on the Pentium M processor.   |
|                  |     | 19                         | Reserved  |
|                  |     | 21:20                      | Symmetric Arbitration ID (R/O)<br>Always 00B on the Pentium M processor.  |
|                  |     | 26:22                      | Clock Frequency Ratio (R/O)   |
| 40H              | 64  | MSR_LASTBRANCH_0           | Last Branch Record 0 (R/W)<br>One of 8 last branch record registers on the last branch record stack: bits 31-0 hold the 'from' address and bits 63-32 hold the to address.<br>See also: <ul style="list-style-type: none"> <li>▪ Last Branch Record Stack TOS at 1C9H.</li> <li>▪ Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)".</li> </ul> |
| 41H              | 65  | MSR_LASTBRANCH_1           | Last Branch Record 1 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 42H              | 66  | MSR_LASTBRANCH_2           | Last Branch Record 2 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 43H              | 67  | MSR_LASTBRANCH_3           | Last Branch Record 3 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 44H              | 68  | MSR_LASTBRANCH_4           | Last Branch Record 4 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 45H              | 69  | MSR_LASTBRANCH_5           | Last Branch Record 5 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 46H              | 70  | MSR_LASTBRANCH_6           | Last Branch Record 6 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 47H              | 71  | MSR_LASTBRANCH_7           | Last Branch Record 7 (R/W)<br>See description of MSR_LASTBRANCH_0.  |
| 119H             | 281 | MSR_BBL_CR_CTL             | Control Register<br>Used to program L2 commands to be issued via cache configuration accesses mechanism. Also receives L2 lookup response.  |
|                  |     | 63:0                       | Reserved  |
| 11EH             | 281 | MSR_BBL_CR_CTL3            | Control register 3<br>Used to configure the L2 Cache.   |

Table 2-49. MSRs in Pentium M Processors (Contd.)

| Register Address |     | Register Name / Bit Fields | Bit Description   |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
|                  |     | 0                          | L2 Hardware Enabled (RO)<br>1 = If the L2 is hardware-enabled.<br>0 = Indicates if the L2 is hardware-disabled.   |
|                  |     | 4:1                        | Reserved  |
|                  |     | 5                          | ECC Check Enable (RO)<br>This bit enables ECC checking on the cache data bus. ECC is always generated on write cycles.<br>0 = Disabled (default)<br>1 = Enabled<br>For the Pentium M processor, ECC checking on the cache data bus is always enabled.           |
|                  |     | 7:6                        | Reserved  |
|                  |     | 8                          | L2 Enabled (R/W)<br>1 = L2 cache has been initialized<br>0 = Disabled (default)<br>Until this bit is set the processor will not respond to the WBINVD instruction or the assertion of the FLUSH# input.   |
|                  |     | 22:9                       | Reserved  |
|                  |     | 23                         | L2 Not Present (RO)<br>0 = L2 Present<br>1 = L2 Not Present   |
|                  |     | 63:24                      | Reserved  |
| 179H             | 377 | IA32_MCG_CAP               | Read-only register that provides information about the machine-check architecture of the processor.   |
|                  |     | 7:0                        | Count (RO)<br>Indicates the number of hardware unit error reporting banks available in the processor.   |
|                  |     | 8                          | IA32_MCG_CTL Present (RO)<br>1 = Indicates that the processor implements the MSR_MCG_CTL register found at MSR 17BH.<br>0 = Not supported.  |
|                  |     | 63:9                       | Reserved  |
| 17AH             | 378 | IA32_MCG_STATUS            | Global Machine Check Status   |
|                  |     | 0                          | RIPV<br>When set, this bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) can be used to restart the program. If this bit is cleared, the program cannot be reliably restarted. |
|                  |     | 1                          | EIPV<br>When set, this bit indicates that the instruction addressed by the instruction pointer pushed on the stack (when the machine check was generated) is directly associated with the error.  |

**Table 2-49. MSRs in Pentium M Processors (Contd.)**

| Register Address |     | Register Name / Bit Fields | Bit Description   |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
|                  |     | 2                          | MCIP<br>When set, this bit indicates that a machine check has been generated. If a second machine check is detected while this bit is still set, the processor enters a shutdown state. Software should write this bit to 0 after processing a machine check exception.   |
|                  |     | 63:3                       | Reserved  |
| 198H             | 408 | IA32_PERF_STATUS           | See Table 2-2.  |
| 199H             | 409 | IA32_PERF_CTL              | See Table 2-2.  |
| 19AH             | 410 | IA32_CLOCK_MODULATION      | Clock Modulation (R/W).<br>See Table 2-2.<br>See Section 14.7.3, "Software Controlled Clock Modulation."  |
| 19BH             | 411 | IA32_THERM_INTERRUPT       | Thermal Interrupt Control (R/W)<br>See Table 2-2.<br>See Section 14.7.2, "Thermal Monitor."   |
| 19CH             | 412 | IA32_THERM_STATUS          | Thermal Monitor Status (R/W)<br>See Table 2-2.<br>See Section 14.7.2, "Thermal Monitor."  |
| 19DH             | 413 | MSR_THERM2_CTL             | Thermal Monitor 2 Control   |
|                  |     | 15:0                       | Reserved  |
|                  |     | 16                         | TM_SELECT (R/W)<br>Mode of automatic thermal monitor:<br>0 = Thermal Monitor 1 (thermally-initiated on-die modulation of the stop-clock duty cycle)<br>1 = Thermal Monitor 2 (thermally-initiated frequency transitions)<br>If bit 3 of the IA32_MISC_ENABLE register is cleared, TM_SELECT has no effect. Neither TM1 nor TM2 will be enabled. |
|                  |     | 63:16                      | Reserved  |
| 1A0H             | 416 | IA32_MISC_ENABLE           | Enable Miscellaneous Processor Features (R/W)<br>Allows a variety of processor functions to be enabled and disabled.  |
|                  |     | 2:0                        | Reserved  |

Table 2-49. MSRs in Pentium M Processors (Contd.)

| Register Address |     | Register Name / Bit Fields | Bit Description   |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
|                  |     | 3                          | <p>Automatic Thermal Control Circuit Enable (R/W)</p> <p>1 = Setting this bit enables the thermal control circuit (TCC) portion of the Intel Thermal Monitor feature. This allows processor clocks to be automatically modulated based on the processor's thermal sensor operation.</p> <p>0 = Disabled (default).</p> <p>The automatic thermal control circuit enable bit determines if the thermal control circuit (TCC) will be activated when the processor's internal thermal sensor determines the processor is about to exceed its maximum operating temperature.</p> <p>When the TCC is activated and TM1 is enabled, the processors clocks will be forced to a 50% duty cycle. BIOS must enable this feature.</p> <p>The bit should not be confused with the on-demand thermal control circuit enable bit.</p> |
|                  |     | 6:4                        | Reserved  |
|                  |     | 7                          | <p>Performance Monitoring Available (R)</p> <p>1 = Performance monitoring enabled.</p> <p>0 = Performance monitoring disabled.</p>  |
|                  |     | 9:8                        | Reserved  |
|                  |     | 10                         | <p>FERR# Multiplexing Enable (R/W)</p> <p>1 = FERR# asserted by the processor to indicate a pending break event within the processor.</p> <p>0 = Indicates compatible FERR# signaling behavior.</p> <p>This bit must be set to 1 to support XAPIC interrupt model usage.</p>  |
|                  |     |                            | <p>Branch Trace Storage Unavailable (RO)</p> <p>1 = Processor doesn't support branch trace storage (BTS)</p> <p>0 = BTS is supported</p>  |
|                  |     | 12                         | <p>Processor Event Based Sampling Unavailable (RO)</p> <p>1 = Processor does not support processor event based sampling (PEBS);</p> <p>0 = PEBS is supported.</p> <p>The Pentium M processor does not support PEBS.</p>   |
|                  |     | 15:13                      | Reserved  |
|                  |     | 16                         | <p>Enhanced Intel SpeedStep Technology Enable (R/W)</p> <p>1 = Enhanced Intel SpeedStep Technology enabled.</p> <p>On the Pentium M processor, this bit may be configured to be read-only.</p>  |
|                  |     | 22:17                      | Reserved  |
|                  |     | 23                         | <p>xTPR Message Disable (R/W)</p> <p>When set to 1, xTPR messages are disabled. xTPR messages are optional messages that allow the processor to inform the chipset of its priority. The default is processor specific.</p>  |
|                  |     | 63:24                      | Reserved  |

**Table 2-49. MSRs in Pentium M Processors (Contd.)**

| Register Address |      | Register Name / Bit Fields | Bit Description   |
|------------------|------|----------------------------|---|
| Hex              | Dec  |                            |   |
| 1C9H             | 457  | MSR_LASTBRANCH_TOS         | Last Branch Record Stack TOS (R/W)<br>Contains an index (bits 0-3) that points to the MSR containing the most recent branch record. See also:<br><ul style="list-style-type: none"> <li>▪ MSR_LASTBRANCH_0_FROM_IP (at 40H).</li> <li>▪ Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)".</li> </ul>   |
| 1D9H             | 473  | MSR_DEBUGCTLB              | Debug Control (R/W)<br>Controls how several debug features are used. Bit definitions are discussed in the referenced section.<br>See Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)."   |
| 1DDH             | 477  | MSR_LER_TO_LIP             | Last Exception Record To Linear IP (R)<br>This area contains a pointer to the target of the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.<br>See Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)" and Section 17.16.2, "Last Branch and Last Exception MSRs." |
| 1DEH             | 478  | MSR_LER_FROM_LIP           | Last Exception Record From Linear IP (R)<br>Contains a pointer to the last branch instruction that the processor executed prior to the last exception that was generated or the last interrupt that was handled.<br>See Section 17.15, "Last Branch, Interrupt, and Exception Recording (Pentium M Processors)" and Section 17.16.2, "Last Branch and Last Exception MSRs."                       |
| 2FFH             | 767  | IA32_MTRR_DEF_TYPE         | Default Memory Types (R/W)<br>Sets the memory type for the regions of physical memory that are not mapped by the MTRRs.<br>See Section 11.11.2.1, "IA32_MTRR_DEF_TYPE MSR."   |
| 400H             | 1024 | IA32_MCO_CTL               | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 401H             | 1025 | IA32_MCO_STATUS            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 402H             | 1026 | IA32_MCO_ADDR              | See Section 14.3.2.3, "IA32_MCi_ADDR MSRs".<br>The IA32_MCO_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MCO_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.  |
| 404H             | 1028 | IA32_MC1_CTL               | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 405H             | 1029 | IA32_MC1_STATUS            | See Section 15.3.2.2, "IA32_MCi_STATUS MSRS."   |
| 406H             | 1030 | IA32_MC1_ADDR              | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs".<br>The IA32_MC1_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC1_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.  |
| 408H             | 1032 | IA32_MC2_CTL               | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."  |
| 409H             | 1033 | IA32_MC2_STATUS            | See Chapter 15.3.2.2, "IA32_MCi_STATUS MSRS."   |

Table 2-49. MSRs in Pentium M Processors (Contd.)

| Register Address |      | Register Name / Bit Fields | Bit Description  |
|------------------|------|----------------------------|--|
| Hex              | Dec  |                            |  |
| 40AH             | 1034 | IA32_MC2_ADDR              | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The IA32_MC2_ADDR register is either not implemented or contains no address if the ADDR_V flag in the IA32_MC2_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception. |
| 40CH             | 1036 | MSR_MC4_CTL                | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 40DH             | 1037 | MSR_MC4_STATUS             | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 40EH             | 1038 | MSR_MC4_ADDR               | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC4_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC4_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 410H             | 1040 | MSR_MC3_CTL                | See Section 15.3.2.1, "IA32_MCi_CTL MSRs."   |
| 411H             | 1041 | MSR_MC3_STATUS             | See Section 15.3.2.2, "IA32_MCi_STATUS MSRs."  |
| 412H             | 1042 | MSR_MC3_ADDR               | See Section 15.3.2.3, "IA32_MCi_ADDR MSRs."<br>The MSR_MC3_ADDR register is either not implemented or contains no address if the ADDR_V flag in the MSR_MC3_STATUS register is clear. When not implemented in the processor, all reads and writes to this MSR will cause a general-protection exception.   |
| 600H             | 1536 | IA32_DS_AREA               | DS Save Area (R/W)<br>See Table 2-2.<br>Points to the DS buffer management area, which is used to manage the BTS and PEBS buffers. See Section 18.6.3.4, "Debug Store (DS) Mechanism."   |
|                  |      | 31:0                       | DS Buffer Management Area<br>Linear address of the first byte of the DS buffer management area.  |
|                  |      | 63:32                      | Reserved   |

## 2.21 MSRS IN THE P6 FAMILY PROCESSORS

The following MSRs are defined for the P6 family processors. The MSRs in this table that are shaded are available only in the Pentium II and Pentium III processors. Beginning with the Pentium 4 processor, some of the MSRs in this list have been designated as "architectural" and have had their names changed. See Table 2-2 for a list of the architectural MSRs.

Table 2-50. MSRs in the P6 Family Processors

| Register Address |     | Register Name / Bit Fields | Bit Description                                 |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
| 0H               | 0   | P5_MC_ADDR                 | See Section 2.22, "MSRs in Pentium Processors." |
| 1H               | 1   | P5_MC_TYPE                 | See Section 2.22, "MSRs in Pentium Processors." |
| 10H              | 16  | TSC                        | See Section 17.17, "Time-Stamp Counter."        |

**Table 2-50. MSRs in the P6 Family Processors (Contd.)**

| Register Address |  | Register Name / Bit Fields | Bit Description   |
|------------------|--|----------------------------|---|
| Hex              | Dec  |                            |   |
| 17H              | 23   | IA32_PLATFORM_ID           | Platform ID (R)<br>The operating system can use this MSR to determine “slot” information for the processor and the proper microcode update to load.   |
|                  |  | 49:0                       | Reserved  |
|                  |  | 52:50                      | Platform Id (R)<br>Contains information concerning the intended platform for the processor.<br>52 51 50<br>0 0 0 Processor Flag 0<br>0 0 1 Processor Flag 1<br>0 1 0 Processor Flag 2<br>0 1 1 Processor Flag 3<br>1 0 0 Processor Flag 4<br>1 0 1 Processor Flag 5<br>1 1 0 Processor Flag 6<br>1 1 1 Processor Flag 7 |
|                  |  | 56:53                      | L2 Cache Latency Read.  |
|                  |  | 59:57                      | Reserved  |
|                  |  | 60                         | Clock Frequency Ratio Read.   |
|                  |  | 63:61                      | Reserved  |
|                  |  | 1BH                        | 27  |
| 7:0              | Reserved   |                            |   |
| 8                | Boot Strap Processor Indicator Bit<br>1 = BSP                                |                            |   |
| 10:9             | Reserved   |                            |   |
| 11               | APIC Global Enable Bit - Permanent till reset<br>1 = Enabled<br>0 = Disabled |                            |   |
| 31:12            | APIC Base Address.   |                            |   |
| 63:32            | Reserved   |                            |   |
| 2AH              | 42   | EBL_CR_POWERON             | Processor Hard Power-On Configuration (R/W) Enables and disables processor features; (R) indicates current processor configuration.   |
|                  |  | 0                          | Reserved <sup>1</sup>   |
|                  |  | 1                          | Data Error Checking Enable (R/W)<br>1 = Enabled<br>0 = Disabled   |
|                  |  | 2                          | Response Error Checking Enable FRCERR Observation Enable (R/W)<br>1 = Enabled<br>0 = Disabled   |
|                  |  | 3                          | AERR# Drive Enable (R/W)<br>1 = Enabled<br>0 = Disabled   |



Table 2-50. MSRs in the P6 Family Processors (Contd.)

| Register Address |     | Register Name / Bit Fields | Bit Description  |
|------------------|-----|----------------------------|--|
| Hex              | Dec |                            |  |
|                  |     | 4                          | BERR# Enable for Initiator Bus Requests (R/W)<br>1 = Enabled<br>0 = Disabled           |
|                  |     | 5                          | Reserved   |
|                  |     | 6                          | BERR# Driver Enable for Initiator Internal Errors (R/W)<br>1 = Enabled<br>0 = Disabled |
|                  |     | 7                          | BINIT# Driver Enable (R/W)<br>1 = Enabled<br>0 = Disabled                              |
|                  |     | 8                          | Output Tri-state Enabled (R)<br>1 = Enabled<br>0 = Disabled                            |
|                  |     | 9                          | Execute BIST (R)<br>1 = Enabled<br>0 = Disabled  |
|                  |     | 10                         | AERR# Observation Enabled (R)<br>1 = Enabled<br>0 = Disabled                           |
|                  |     | 11                         | Reserved   |
|                  |     | 12                         | BINIT# Observation Enabled (R)<br>1 = Enabled<br>0 = Disabled                          |
|                  |     | 13                         | In Order Queue Depth (R)<br>1 = 1<br>0 = 8   |
|                  |     | 14                         | 1-MByte Power on Reset Vector (R)<br>1 = 1MByte<br>0 = 4GBytes                         |
|                  |     | 15                         | FRC Mode Enable (R)<br>1 = Enabled<br>0 = Disabled                                     |
|                  |     | 17:16                      | APIC Cluster ID (R)  |
|                  |     | 19:18                      | System Bus Frequency (R)<br>00 = 66MHz<br>10 = 100MHz<br>01 = 133MHz<br>11 = Reserved  |
|                  |     | 21:20                      | Symmetric Arbitration ID (R)   |
|                  |     | 25:22                      | Clock Frequency Ratio (R)  |
|                  |     | 26                         | Low Power Mode Enable (R/W)  |

**Table 2-50. MSRs in the P6 Family Processors (Contd.)**

| Register Address |     | Register Name / Bit Fields | Bit Description   |
|------------------|-----|----------------------------|---|
| Hex              | Dec |                            |   |
|                  |     | 27                         | Clock Frequency Ratio   |
|                  |     | 63:28                      | Reserved <sup>1</sup>   |
| 33H              | 51  | TEST_CTL                   | Test Control Register   |
|                  |     | 29:0                       | Reserved  |
|                  |     | 30                         | Streaming Buffer Disable  |
|                  |     | 31                         | Disable LOCK#<br>Assertion for split locked access.   |
| 79H              | 121 | BIOS_UPDT_TRIG             | BIOS Update Trigger Register.   |
| 88H              | 136 | BBL_CR_D0[63:0]            | Chunk 0 data register D[63:0]; used to write to and read from the L2  |
| 89H              | 137 | BBL_CR_D1[63:0]            | Chunk 1 data register D[63:0]; used to write to and read from the L2  |
| 8AH              | 138 | BBL_CR_D2[63:0]            | Chunk 2 data register D[63:0]; used to write to and read from the L2  |
| 8BH              | 139 | BIOS_SIGN/BBL_CR_D3[63:0]  | BIOS Update Signature Register or Chunk 3 data register D[63:0]<br>Used to write to and read from the L2 depending on the usage model.  |
| C1H              | 193 | PerfCtr0 (PERFCTR0)        | Performance Counter Register<br>See Table 2-2.  |
| C2H              | 194 | PerfCtr1 (PERFCTR1)        | Performance Counter Register<br>See Table 2-2.  |
| FEH              | 254 | MTRRcap                    | Memory Type Range Registers   |
| 116H             | 278 | BBL_CR_ADDR [63:0]         | Address register: used to send specified address (A31-A3) to L2 during cache initialization accesses.                                   |
|                  |     | BBL_CR_ADDR [63:32]        | Reserved,   |
|                  |     | BBL_CR_ADDR [31:3]         | Address bits [35:3]   |
|                  |     | BBL_CR_ADDR [2:0]          | Reserved Set to 0.  |
| 118H             | 280 | BBL_CR_DECC[63:0]          | Data ECC register D[7:0]: used to write ECC and read ECC to/from L2   |
| 119H             | 281 | BBL_CR_CTL                 | Control register: used to program L2 commands to be issued via cache configuration accesses mechanism. Also receives L2 lookup response |
|                  |     | BL_CR_CTL[63:22]           | Reserved  |
|                  |     | BBL_CR_CTL[21]             | Processor number <sup>2</sup><br>Disable = 1<br>Enable = 0<br>Reserved  |

Table 2-50. MSRs in the P6 Family Processors (Contd.)

| Register Address |     | Register Name / Bit Fields   | Bit Description   |
|------------------|-----|--|---|
| Hex              | Dec |  |   |
|                  |     | BBL_CR_CTL[20:19]<br>BBL_CR_CTL[18]<br>BBL_CR_CTL[17]<br>BBL_CR_CTL[16]<br>BBL_CR_CTL[15:14]<br>BBL_CR_CTL[13:12]<br><br>BBL_CR_CTL[11:10]<br><br>BBL_CR_CTL[9:8]<br>BBL_CR_CTL[7]<br>BBL_CR_CTL[6:5]                      | User supplied ECC<br>Reserved<br>L2 Hit<br>Reserved<br>State from L2<br>Modified - 11, Exclusive - 10, Shared - 01, Invalid - 00<br>Way from L2<br>Way 0 - 00, Way 1 - 01, Way 2 - 10, Way 3 - 11<br>Way to L2<br>Reserved<br>State to L2   |
|                  |     | BBL_CR_CTL[4:0]<br>01100<br>01110<br>01111<br>00010<br>00011<br>010 + MESI encode<br>111 + MESI encode<br>100 + MESI encode  | L2 Command<br>Data Read w/ LRU update (RLU)<br>Tag Read w/ Data Read (TRR)<br>Tag Inquire (TI)<br>L2 Control Register Read (CR)<br>L2 Control Register Write (CW)<br>Tag Write w/ Data Read (TWR)<br>Tag Write w/ Data Write (TWW)<br>Tag Write (TW)  |
| 11AH             | 282 | BBL_CR_TRIG  | Trigger register: used to initiate a cache configuration accesses access, Write only with Data = 0.   |
| 11BH             | 283 | BBL_CR_BUSY  | Busy register: indicates when a cache configuration accesses L2 command is in progress. D[0] = 1 = BUSY   |
| 11EH             | 286 | BBL_CR_CTL3<br><br>BBL_CR_CTL3[63:26]<br>BBL_CR_CTL3[25]<br>BBL_CR_CTL3[24]<br>BBL_CR_CTL3[23]<br><br>BBL_CR_CTL3[22:20]<br>111<br>110<br>101<br>100<br>011<br>010<br>001<br>000<br><br>BBL_CR_CTL3[19]<br>BBL_CR_CTL3[18] | Control register 3: used to configure the L2 Cache<br><br>Reserved<br>Cache bus fraction (read only)<br>Reserved<br>L2 Hardware Disable (read only)<br><br>L2 Physical Address Range support<br>64GBytes<br>32GBytes<br>16GBytes<br>8GBytes<br>4GBytes<br>2GBytes<br>1GBytes<br>512MBytes<br><br>Reserved<br>Cache State error checking enable (read/write) |

**Table 2-50. MSRs in the P6 Family Processors (Contd.)**

| Register Address |     | Register Name / Bit Fields   | Bit Description  |
|------------------|-----|--|--|
| Hex              | Dec |  |  |
|                  |     | BBL_CR_CTL3[17:13]<br>00001<br>00010<br>00100<br>01000<br>10000<br><br>BBL_CR_CTL3[12:11]<br>BBL_CR_CTL3[10:9]<br>00<br>01<br>10<br>11<br><br>BBL_CR_CTL3[8]<br>BBL_CR_CTL3[7]<br>BBL_CR_CTL3[6]<br>BBL_CR_CTL3[5]<br>BBL_CR_CTL3[4:1]<br>BBL_CR_CTL3[0] | Cache size per bank (read/write)<br>256KBytes<br>512KBytes<br>1MByte<br>2MByte<br>4MBytes<br><br>Number of L2 banks (read only)<br>L2 Associativity (read only)<br>Direct Mapped<br>2 Way<br>4 Way<br>Reserved<br><br>L2 Enabled (read/write)<br>CRTN Parity Check Enable (read/write)<br>Address Parity Check Enable (read/write)<br>ECC Check Enable (read/write)<br>L2 Cache Latency (read/write)<br>L2 Configured (read/write<br>) |
| 174H             | 372 | SYSENTER_CS_MSR  | CS register target for CPL 0 code  |
| 175H             | 373 | SYSENTER_ESP_MSR   | Stack pointer for CPL 0 stack  |
| 176H             | 374 | SYSENTER_EIP_MSR   | CPL 0 code entry point   |
| 179H             | 377 | MCG_CAP  | Machine Check Global Control Register  |
| 17AH             | 378 | MCG_STATUS   | Machine Check Error Reporting Register - contains information related to machine-check error if its VAL (valid) flag is set. Software is responsible for clearing IA32_MCi_STATUS MSRs by explicitly writing 0s to them; writing 1s to them causes a general-protection exception.   |
| 17BH             | 379 | MCG_CTL  | Machine Check Error Reporting Register - controls signaling of #MC for errors produced by a particular hardware unit (or group of hardware units).   |
| 186H             | 390 | PerfEvtSel0 (EVNTSEL0)   | Performance Event Select Register 0 (R/W)  |
|                  |     | 7:0  | Event Select<br>Refer to Performance Counter section for a list of event encodings.  |
|                  |     | 15:8   | UMASK (Unit Mask)<br>Unit mask register set to 0 to enable all count options.  |
|                  |     | 16   | USER<br>Controls the counting of events at Privilege levels of 1, 2, and 3.  |
|                  |     | 17   | OS<br>Controls the counting of events at Privilege level of 0.   |

Table 2-50. MSRs in the P6 Family Processors (Contd.)

| Register Address |     | Register Name / Bit Fields | Bit Description  |
|------------------|-----|----------------------------|--|
| Hex              | Dec |                            |  |
|                  |     | 18                         | E<br>Occurrence/Duration Mode Select<br>1 = Occurrence<br>0 = Duration                             |
|                  |     | 19                         | PC<br>Enabled the signaling of performance counter overflow via BPO pin                            |
|                  |     | 20                         | INT<br>Enables the signaling of counter overflow via input to APIC<br>1 = Enable<br>0 = Disable    |
|                  |     | 22                         | ENABLE<br>Enables the counting of performance events in both counters<br>1 = Enable<br>0 = Disable |
|                  |     | 23                         | INV<br>Inverts the result of the CMASK condition<br>1 = Inverted<br>0 = Non-Inverted               |
|                  |     | 31:24                      | CMASK (Counter Mask)   |
| 187H             | 391 | PerfEvtSel1 (EVNTSEL1)     | Performance Event Select for Counter 1 (R/W)   |
|                  |     | 7:0                        | Event Select<br>Refer to Performance Counter section for a list of event encodings.                |
|                  |     | 15:8                       | UMASK (Unit Mask)<br>Unit mask register set to 0 to enable all count options.                      |
|                  |     | 16                         | USER<br>Controls the counting of events at Privilege levels of 1, 2, and 3.                        |
|                  |     | 17                         | OS<br>Controls the counting of events at Privilege level of 0.                                     |
|                  |     | 18                         | E<br>Occurrence/Duration Mode Select.<br>1 = Occurrence<br>0 = Duration                            |
|                  |     | 19                         | PC<br>Enabled the signaling of performance counter overflow via BPO pin.                           |

**Table 2-50. MSRs in the P6 Family Processors (Contd.)**

| Register Address |          | Register Name / Bit Fields | Bit Description  |
|------------------|----------|----------------------------|--|
| Hex              | Dec      |                            |  |
|                  |          | 20                         | INT<br>Enables the signaling of counter overflow via input to APIC.<br>1 = Enable<br>0 = Disable   |
|                  |          | 23                         | INV<br>Inverts the result of the CMASK condition.<br>1 = Inverted<br>0 = Non-Inverted  |
|                  |          | 31:24                      | CMASK (Counter Mask)   |
| 1D9H             | 473      | DEBUGCTLMR                 | Enables last branch, interrupt, and exception recording; taken branch breakpoints; the breakpoint reporting pins; and trace messages. This register can be written to using the WRMSR instruction, when operating at privilege level 0 or when in real-address mode. |
|                  |          | 0                          | Enable/Disable Last Branch Records   |
|                  |          | 1                          | Branch Trap Flag   |
|                  |          | 2                          | Performance Monitoring/Break Point Pins  |
|                  |          | 3                          | Performance Monitoring/Break Point Pins  |
|                  |          | 4                          | Performance Monitoring/Break Point Pins  |
|                  |          | 5                          | Performance Monitoring/Break Point Pins  |
|                  |          | 6                          | Enable/Disable Execution Trace Messages  |
| 31:7             | Reserved |                            |  |
| 1DBH             | 475      | LASTBRANCHFROMIP           | 32-bit register for recording the instruction pointers for the last branch, interrupt, or exception that the processor took prior to a debug exception being generated.  |
| 1DCH             | 476      | LASTBRANCHTOIP             | 32-bit register for recording the instruction pointers for the last branch, interrupt, or exception that the processor took prior to a debug exception being generated.  |
| 1DDH             | 477      | LASTINTFROMIP              | Last INT from IP   |
| 1DEH             | 478      | LASTINTTOIP                | Last INT to IP   |
| 200H             | 512      | MTRRphysBase0              | Memory Type Range Registers  |
| 201H             | 513      | MTRRphysMask0              | Memory Type Range Registers  |
| 202H             | 514      | MTRRphysBase1              | Memory Type Range Registers  |
| 203H             | 515      | MTRRphysMask1              | Memory Type Range Registers  |
| 204H             | 516      | MTRRphysBase2              | Memory Type Range Registers  |
| 205H             | 517      | MTRRphysMask2              | Memory Type Range Registers  |
| 206H             | 518      | MTRRphysBase3              | Memory Type Range Registers  |
| 207H             | 519      | MTRRphysMask3              | Memory Type Range Registers  |
| 208H             | 520      | MTRRphysBase4              | Memory Type Range Registers  |
| 209H             | 521      | MTRRphysMask4              | Memory Type Range Registers  |
| 20AH             | 522      | MTRRphysBase5              | Memory Type Range Registers  |

Table 2-50. MSRs in the P6 Family Processors (Contd.)

| Register Address |             | Register Name / Bit Fields | Bit Description  |
|------------------|-------------|----------------------------|--|
| Hex              | Dec         |                            |  |
| 20BH             | 523         | MTRRphysMask5              | Memory Type Range Registers  |
| 20CH             | 524         | MTRRphysBase6              | Memory Type Range Registers  |
| 20DH             | 525         | MTRRphysMask6              | Memory Type Range Registers  |
| 20EH             | 526         | MTRRphysBase7              | Memory Type Range Registers  |
| 20FH             | 527         | MTRRphysMask7              | Memory Type Range Registers  |
| 250H             | 592         | MTRRfix64K_00000           | Memory Type Range Registers  |
| 258H             | 600         | MTRRfix16K_80000           | Memory Type Range Registers  |
| 259H             | 601         | MTRRfix16K_A0000           | Memory Type Range Registers  |
| 268H             | 616         | MTRRfix4K_C0000            | Memory Type Range Registers  |
| 269H             | 617         | MTRRfix4K_C8000            | Memory Type Range Registers  |
| 26AH             | 618         | MTRRfix4K_D0000            | Memory Type Range Registers  |
| 26BH             | 619         | MTRRfix4K_D8000            | Memory Type Range Registers  |
| 26CH             | 620         | MTRRfix4K_E0000            | Memory Type Range Registers  |
| 26DH             | 621         | MTRRfix4K_E8000            | Memory Type Range Registers  |
| 26EH             | 622         | MTRRfix4K_F0000            | Memory Type Range Registers  |
| 26FH             | 623         | MTRRfix4K_F8000            | Memory Type Range Registers  |
| 2FFH             | 767         | MTRRdefType                | Memory Type Range Registers  |
|                  |             | 2:0                        | Default memory type  |
|                  |             | 10                         | Fixed MTRR enable  |
|                  |             | 11                         | MTRR Enable  |
| 400H             | 1024        | MCO_CTL                    | Machine Check Error Reporting Register - controls signaling of #MC for errors produced by a particular hardware unit (or group of hardware units).   |
| 401H             | 1025        | MCO_STATUS                 | Machine Check Error Reporting Register - contains information related to a machine-check error if its VAL (valid) flag is set. Software is responsible for clearing IA32_MCI_STATUS MSRs by explicitly writing 0s to them; writing 1s to them causes a general-protection exception. |
|                  |             | 15:0                       | MC_STATUS_MCACOD   |
|                  |             | 31:16                      | MC_STATUS_MSCOD  |
|                  |             | 57                         | MC_STATUS_DAM  |
|                  |             | 58                         | MC_STATUS_ADDRV  |
|                  |             | 59                         | MC_STATUS_MISCV  |
|                  |             | 60                         | MC_STATUS_EN. (Note: For MCO_STATUS only, this bit is hardcoded to 1.)   |
|                  |             | 61                         | MC_STATUS_UC   |
|                  |             | 62                         | MC_STATUS_O  |
| 63               | MC_STATUS_V |                            |  |
| 402H             | 1026        | MCO_ADDR                   |  |
| 403H             | 1027        | MCO_MISC                   | Defined in MCA architecture but not implemented in the P6 family processors.   |

Table 2-50. MSRs in the P6 Family Processors (Contd.)

| Register Address |      | Register Name / Bit Fields | Bit Description  |
|------------------|------|----------------------------|--|
| Hex              | Dec  |                            |  |
| 404H             | 1028 | MC1_CTL                    |  |
| 405H             | 1029 | MC1_STATUS                 | Bit definitions same as MCO_STATUS.  |
| 406H             | 1030 | MC1_ADDR                   |  |
| 407H             | 1031 | MC1_MISC                   | Defined in MCA architecture but not implemented in the P6 family processors.         |
| 408H             | 1032 | MC2_CTL                    |  |
| 409H             | 1033 | MC2_STATUS                 | Bit definitions same as MCO_STATUS.  |
| 40AH             | 1034 | MC2_ADDR                   |  |
| 40BH             | 1035 | MC2_MISC                   | Defined in MCA architecture but not implemented in the P6 family processors.         |
| 40CH             | 1036 | MC4_CTL                    |  |
| 40DH             | 1037 | MC4_STATUS                 | Bit definitions same as MCO_STATUS, except bits 0, 4, 57, and 61 are hardcoded to 1. |
| 40EH             | 1038 | MC4_ADDR                   | Defined in MCA architecture but not implemented in P6 Family processors.             |
| 40FH             | 1039 | MC4_MISC                   | Defined in MCA architecture but not implemented in the P6 family processors.         |
| 410H             | 1040 | MC3_CTL                    |  |
| 411H             | 1041 | MC3_STATUS                 | Bit definitions same as MCO_STATUS.  |
| 412H             | 1042 | MC3_ADDR                   |  |
| 413H             | 1043 | MC3_MISC                   | Defined in MCA architecture but not implemented in the P6 family processors.         |

**NOTES**

1. Bit 0 of this register has been redefined several times, and is no longer used in P6 family processors.
2. The processor number feature may be disabled by setting bit 21 of the BBL\_CR\_CTL MSR (model-specific register address 119h) to "1". Once set, bit 21 of the BBL\_CR\_CTL may not be cleared. This bit is write-once. The processor number feature will be disabled until the processor is reset.
3. The Pentium III processor will prevent FSB frequency overclocking with a new shutdown mechanism. If the FSB frequency selected is greater than the internal FSB frequency the processor will shutdown. If the FSB selected is less than the internal FSB frequency the BIOS may choose to use bit 11 to implement its own shutdown policy.

## 2.22 MSRS IN PENTIUM PROCESSORS

The following MSRs are defined for the Pentium processors. The P5\_MC\_ADDR, P5\_MC\_TYPE, and TSC MSRs (named IA32\_P5\_MC\_ADDR, IA32\_P5\_MC\_TYPE, and IA32\_TIME\_STAMP\_COUNTER in the Pentium 4 processor) are architectural; that is, code that accesses these registers will run on Pentium 4 and P6 family processors without generating exceptions (see Section 2.1, "Architectural MSRs"). The CESR, CTR0, and CTR1 MSRs are unique to Pentium processors; code that accesses these registers will generate exceptions on Pentium 4 and P6 family processors.



Table 2-51. MSRs in the Pentium Processor

| Register Address |     | Register Name | Bit Description  |
|------------------|-----|---------------|--|
| Hex              | Dec |               |  |
| 0H               | 0   | P5_MC_ADDR    | See Section 15.10.2, "Pentium Processor Machine-Check Exception Handling." |
| 1H               | 1   | P5_MC_TYPE    | See Section 15.10.2, "Pentium Processor Machine-Check Exception Handling." |
| 10H              | 16  | TSC           | See Section 17.17, "Time-Stamp Counter."                                   |
| 11H              | 17  | CESR          | See Section 18.6.9.1, "Control and Event Select Register (CESR)."          |
| 12H              | 18  | CTRO          | Section 18.6.9.3, "Events Counted."  |
| 13H              | 19  | CTR1          | Section 18.6.9.3, "Events Counted."  |

## 2.23 MSR INDEX

MSRs of recent processors are indexed here for convenience. IA32 MSRs are excluded from this index.

| MSR Name and CPUID DisplayFamily_DisplayModel        | Location       |
|--|----------------|
| MSR_ALF_ESCR0  |                |
| 0FH .....  | See Table 2-45 |
| MSR_ALF_ESCR1  |                |
| 0FH .....  | See Table 2-45 |
| MSR_ANY_CORE_C0                                      |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38 |
| MSR_ANY_GFXE_C0                                      |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38 |
| MSR_BO_PMON_BOX_CTRL                                 |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_BOX_OVF_CTRL                             |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_BOX_STATUS                               |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_CTRO                                     |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_CTR1                                     |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_CTR2                                     |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_CTR3                                     |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_EVNT_SELO                                |                |
| 06_2EH .....   | See Table 2-16 |
| MSR_BO_PMON_EVNT_SEL1                                |                |
| 06_2EH .....   | See Table 2-16 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| MSR_BO_PMON_EVNT_SEL2<br>06_2EH .....         | See Table 2-16 |
| MSR_BO_PMON_EVNT_SEL3<br>06_2EH .....         | See Table 2-16 |
| MSR_BO_PMON_MASK<br>06_2EH .....              | See Table 2-16 |
| MSR_BO_PMON_MATCH<br>06_2EH .....             | See Table 2-16 |
| MSR_B1_PMON_BOX_CTRL<br>06_2EH .....          | See Table 2-16 |
| MSR_B1_PMON_BOX_OVF_CTRL<br>06_2EH .....      | See Table 2-16 |
| MSR_B1_PMON_BOX_STATUS<br>06_2EH .....        | See Table 2-16 |
| MSR_B1_PMON_CTRL0<br>06_2EH .....             | See Table 2-16 |
| MSR_B1_PMON_CTRL1<br>06_2EH .....             | See Table 2-16 |
| MSR_B1_PMON_CTRL2<br>06_2EH .....             | See Table 2-16 |
| MSR_B1_PMON_CTRL3<br>06_2EH .....             | See Table 2-16 |
| MSR_B1_PMON_EVNT_SELO<br>06_2EH .....         | See Table 2-16 |
| MSR_B1_PMON_EVNT_SEL1<br>06_2EH .....         | See Table 2-16 |
| MSR_B1_PMON_EVNT_SEL2<br>06_2EH .....         | See Table 2-16 |
| MSR_B1_PMON_EVNT_SEL3<br>06_2EH .....         | See Table 2-16 |
| MSR_B1_PMON_MASK<br>06_2EH .....              | See Table 2-16 |
| MSR_B1_PMON_MATCH<br>06_2EH .....             | See Table 2-16 |
| MSR_BBL_CR_CTL<br>06_09H .....                | See Table 2-49 |
| MSR_BBL_CR_CTL3<br>06_0FH, 06_17H .....       | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....  | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....  | See Table 2-7  |
| 06_0EH .....                                  | See Table 2-48 |
| 06_09H .....                                  | See Table 2-49 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_BPU_CCCR0  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_CCCR1  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_CCCR2  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_CCCR3  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_COUNTER0                                     |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_COUNTER1                                     |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_COUNTER2                                     |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_COUNTER3                                     |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_ESCR0  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BPU_ESCR1  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BR_DETECT_COUNTER_CONFIG_j                       |                 |
| 06_66H.....  | See Table 2-41  |
| MSR_BR_DETECT_CTRL                                   |                 |
| 06_66H.....  | See Table 2-41  |
| MSR_BR_DETECT_STATUS                                 |                 |
| 06_66H.....  | See Table 2-41  |
| MSR_BSU_ESCR0  |                 |
| OFH .....  | See Table 2-45  |
| MSR_BSU_ESCR1  |                 |
| OFH .....  | See Table 2-45  |
| MSR_CO_PMON_BOX_CTRL                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_CO_PMON_BOX_FILTER                               |                 |
| 06_2DH .....   | See Table 2-23  |
| MSR_CO_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_CO_PMON_BOX_FILTER1                              |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_CO_PMON_BOX_OVF_CTRL                             |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_2EH .....                                  | See Table 2-16 |
| MSR_CO_PMON_BOX_STATUS                        |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR0                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR1                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR2                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR3                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR4                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_CO_PMON_CTR5                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_CO_PMON_EVNT_SELO                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_EVNT_SEL1                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR1                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_EVNT_SEL2                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_CO_PMON_CTR2                              |                |
| 06_2EH .....                                  | See Table 2-16 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_CO_PMON_EVNT_SEL3</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_CO_PMON_EVNT_SEL4</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_CO_PMON_EVNT_SEL5</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C1_PMON_BOX_CTRL</b>                          |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_BOX_FILTER</b>                        |                 |
| 06_2DH .....   | See Table 2-23  |
| <b>MSR_C1_PMON_BOX_FILTER0</b>                       |                 |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_BOX_FILTER1</b>                       |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C1_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_CTRL0</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_CTRL1</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_CTRL2</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C1_PMON_CTRL3</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| MSR_C1_PMON_CTR4                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_C1_PMON_CTR5                              |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_C1_PMON_EVNT_SEL0                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C1_PMON_EVNT_SEL1                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C1_PMON_EVNT_SEL2                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C1_PMON_EVNT_SEL3                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| 06_2DH .....                                  | See Table 2-23 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C1_PMON_EVNT_SEL4                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_C1_PMON_EVNT_SEL5                         |                |
| 06_2EH .....                                  | See Table 2-16 |
| MSR_C10_PMON_BOX_FILTER                       |                |
| 06_3EH .....                                  | See Table 2-27 |
| MSR_C10_PMON_BOX_FILTER0                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C10_PMON_BOX_FILTER1                      |                |
| 06_3EH .....                                  | See Table 2-27 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C11_PMON_BOX_FILTER                       |                |
| 06_3EH .....                                  | See Table 2-27 |
| MSR_C11_PMON_BOX_FILTER0                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C11_PMON_BOX_FILTER1                      |                |
| 06_3EH .....                                  | See Table 2-27 |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C12_PMON_BOX_FILTER                       |                |
| 06_3EH .....                                  | See Table 2-27 |
| MSR_C12_PMON_BOX_FILTER0                      |                |
| 06_3FH .....                                  | See Table 2-32 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_C12_PMON_BOX_FILTER1                             |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C13_PMON_BOX_FILTER                              |                 |
| 06_3EH .....   | See Table 2-27  |
| MSR_C13_PMON_BOX_FILTER0                             |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C13_PMON_BOX_FILTER1                             |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C14_PMON_BOX_FILTER                              |                 |
| 06_3EH .....   | See Table 2-27  |
| MSR_C14_PMON_BOX_FILTER0                             |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C14_PMON_BOX_FILTER1                             |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_BOX_CTL                                 |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_BOX_FILTER0                             |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_BOX_FILTER1                             |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_BOX_STATUS                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_CTR0                                    |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_CTR1                                    |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_CTR2                                    |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_CTR3                                    |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_EVNTSELO                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_EVNTSEL1                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_EVNTSEL2                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C15_PMON_EVNTSEL3                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C16_PMON_BOX_CTL                                 |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_BOX_FILTER0                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_BOX_FILTER1                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_BOX_STATUS                       |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_CTR0                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_CTR3                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_CTR2                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_CTR3                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_EVNTSELO                         |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_EVNTSEL1                         |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_EVNTSEL2                         |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C16_PMON_EVNTSEL3                         |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_BOX_CTL                          |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_BOX_FILTER0                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_BOX_FILTER1                      |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_BOX_STATUS                       |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_CTR0                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_CTR1                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_CTR2                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_CTR3                             |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_EVNTSELO                         |                |
| 06_3FH .....                                  | See Table 2-32 |
| MSR_C17_PMON_EVNTSEL1                         |                |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C17_PMON_EVNTSEL2                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C17_PMON_EVNTSEL3                                |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_BOX_CTRL                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_BOX_FILTER                               |                 |
| 06_2DH .....   | See Table 2-23  |
| MSR_C2_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_BOX_FILTER1                              |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_BOX_OVF_CTRL                             |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C2_PMON_BOX_STATUS                               |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_CTRL0                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_CTRL1                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_CTRL2                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_CTRL3                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C2_PMON_CTRL4                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C2_PMON_CTRL5                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C2_PMON_EVNT_SELO                                |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C2_PMON_EVNT_SEL1</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C2_PMON_EVNT_SEL2</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C2_PMON_EVNT_SEL3</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C2_PMON_EVNT_SEL4</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C2_PMON_EVNT_SEL5</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C3_PMON_BOX_CTRL</b>                          |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C3_PMON_BOX_FILTER</b>                        |                 |
| 06_2DH .....   | See Table 2-23  |
| <b>MSR_C3_PMON_BOX_FILTER0</b>                       |                 |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C3_PMON_BOX_FILTER1</b>                       |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C3_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C3_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C3_PMON_CTRL0</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C3_PMON_CTRL1</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_CTR2                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_CTR3                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_CTR4                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C3_PMON_CTR5                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C3_PMON_EVNT_SEL0                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_EVNT_SEL1                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_EVNT_SEL2                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_EVNT_SEL3                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C3_PMON_EVNT_SEL4                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C3_PMON_EVNT_SEL5                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C4_PMON_BOX_CTRL                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C4_PMON_BOX_FILTER                               |                 |
| 06_2DH .....   | See Table 2-23  |
| MSR_C4_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C4_PMON_BOX_FILTER1                              |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C4_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_CTRL0</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_CTRL1</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_CTRL2</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_CTRL3</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_CTRL4</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C4_PMON_CTRL5</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C4_PMON_EVNT_SELO</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_EVNT_SEL1</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_EVNT_SEL2</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C4_PMON_EVNT_SEL3</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C4_PMON_EVNT_SEL4                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C4_PMON_EVNT_SEL5                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C5_PMON_BOX_CTRL                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_BOX_FILTER                               |                 |
| 06_2DH .....   | See Table 2-23  |
| MSR_C5_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_BOX_FILTER1                              |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_BOX_OVF_CTRL                             |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C5_PMON_BOX_STATUS                               |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_CTRL0                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_CTRL1                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_CTRL2                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_CTRL3                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C5_PMON_CTRL4                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C5_PMON_CTRL5                                    |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C5_PMON_EVNT_SEL0                                |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C5_PMON_EVNT_SEL1</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C5_PMON_EVNT_SEL2</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C5_PMON_EVNT_SEL3</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C5_PMON_EVNT_SEL4</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C5_PMON_EVNT_SEL5</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C6_PMON_BOX_CTRL</b>                          |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C6_PMON_BOX_FILTER</b>                        |                 |
| 06_2DH .....   | See Table 2-23  |
| <b>MSR_C6_PMON_BOX_FILTER0</b>                       |                 |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C6_PMON_BOX_FILTER1</b>                       |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C6_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C6_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C6_PMON_CTRL0</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C6_PMON_CTRL1</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_CTR2                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_CTR3                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_CTR4                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C6_PMON_CTR5                                     |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C6_PMON_EVNT_SEL0                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_EVNT_SEL1                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_EVNT_SEL2                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_EVNT_SEL3                                |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C6_PMON_EVNT_SEL4                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C6_PMON_EVNT_SEL5                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C7_PMON_BOX_CTRL                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C7_PMON_BOX_FILTER                               |                 |
| 06_2DH .....   | See Table 2-23  |
| MSR_C7_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C7_PMON_BOX_FILTER1                              |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C7_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_CTRL0</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_CTRL1</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_CTRL2</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_CTRL3</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_CTRL4</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C7_PMON_CTRL5</b>                             |                 |
| 06_2EH .....   | See Table 2-16  |
| <b>MSR_C7_PMON_EVNT_SELO</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_EVNT_SEL1</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_EVNT_SEL2</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C7_PMON_EVNT_SEL3</b>                         |                 |
| 06_2EH .....   | See Table 2-16  |
| 06_2DH .....   | See Table 2-23  |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C7_PMON_EVNT_SEL4                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C7_PMON_EVNT_SEL5                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_C8_PMON_BOX_CTRL                                 |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_BOX_FILTER                               |                 |
| 06_3EH .....   | See Table 2-27  |
| MSR_C8_PMON_BOX_FILTER0                              |                 |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_BOX_FILTER1                              |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_BOX_OVF_CTRL                             |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C8_PMON_BOX_STATUS                               |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_CTRL0                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_CTRL1                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_CTRL2                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_CTRL3                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C8_PMON_CTRL4                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C8_PMON_CTRL5                                    |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C8_PMON_EVNT_SEL0                                |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C8_PMON_EVNT_SEL1</b>                         |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C8_PMON_EVNT_SEL2</b>                         |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C8_PMON_EVNT_SEL3</b>                         |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C8_PMON_EVNT_SEL4</b>                         |                 |
| 06_2FH .....   | See Table 2-18  |
| <b>MSR_C8_PMON_EVNT_SEL5</b>                         |                 |
| 06_2FH .....   | See Table 2-18  |
| <b>MSR_C9_PMON_BOX_CTRL</b>                          |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C9_PMON_BOX_FILTER</b>                        |                 |
| 06_3EH .....   | See Table 2-27  |
| <b>MSR_C9_PMON_BOX_FILTER0</b>                       |                 |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C9_PMON_BOX_FILTER1</b>                       |                 |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C9_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2FH .....   | See Table 2-18  |
| <b>MSR_C9_PMON_BOX_STATUS</b>                        |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C9_PMON_CTRL0</b>                             |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| <b>MSR_C9_PMON_CTRL1</b>                             |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_CTR2                                     |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_CTR3                                     |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_CTR4                                     |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C9_PMON_CTR5                                     |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C9_PMON_EVNT_SEL0                                |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_EVNT_SEL1                                |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_EVNT_SEL2                                |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_EVNT_SEL3                                |                 |
| 06_2FH .....   | See Table 2-18  |
| 06_3EH .....   | See Table 2-27  |
| 06_3FH .....   | See Table 2-32  |
| MSR_C9_PMON_EVNT_SEL4                                |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_C9_PMON_EVNT_SEL5                                |                 |
| 06_2FH .....   | See Table 2-18  |
| MSR_CC6_DEMOTION_POLICY_CONFIG                       |                 |
| 06_37H .....   | See Table 2-9   |
| MSR_CONFIG_TDP_CONTROL                               |                 |
| 06_3AH .....   | See Table 2-24  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-28  |
| 06_57H .....   | See Table 2-43  |
| MSR_CONFIG_TDP_LEVEL1                                |                 |
| 06_3AH .....   | See Table 2-24  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-28  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                 | <b>Location</b> |
|--|-----------------|
| 06_57H .....   | See Table 2-43  |
| <b>MSR_CONFIG_TDP_LEVEL2</b>   |                 |
| 06_3AH .....   | See Table 2-24  |
| 06_3CH, 06_45H, 06_46H .....   | See Table 2-28  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_CONFIG_TDP_NOMINAL</b>  |                 |
| 06_3AH .....   | See Table 2-24  |
| 06_3CH, 06_45H, 06_46H .....   | See Table 2-28  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_CORE_C1_RESIDENCY</b>   |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH .....         | See Table 2-6   |
| 06_66H .....   | See Table 2-41  |
| <b>MSR_CORE_C3_RESIDENCY</b>   |                 |
| 06_5CH, 06_7AH .....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH, 06_2FH .....         | See Table 2-14  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H ..... | See Table 2-19  |
| <b>MSR_CORE_C6_RESIDENCY</b>   |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH .....         | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH, 06_2FH .....         | See Table 2-14  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H ..... | See Table 2-19  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_CORE_C7_RESIDENCY</b>   |                 |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H ..... | See Table 2-19  |
| <b>MSR_CORE_GFXE_OVERLAP_CO</b>                                      |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H .....                 | See Table 2-38  |
| <b>MSR_CORE_HDC_RESIDENCY</b>  |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H .....                 | See Table 2-38  |
| <b>MSR_CORE_PERF_LIMIT_REASONS</b>                                   |                 |
| 06_5CH, 06_7AH .....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H .....   | See Table 2-29  |
| 06_3F .....  | See Table 2-31  |
| 06_56H, 06_4FH .....   | See Table 2-35  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_CORE_THREAD_COUNT</b>   |                 |
| 06_3FH .....   | See Table 2-31  |
| <b>MSR_CRU_ESCR0</b>   |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_CRU_ESCR1</b>   |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_CRU_ESCR2</b>   |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_CRU_ESCR3</b>   |                 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 0FH .....  | See Table 2-45  |
| <b>MSR_CRU_ESCR4</b>                                 |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_CRU_ESCR5</b>                                 |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_DAC_ESCR0</b>                                 |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_DAC_ESCR1</b>                                 |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_DRAM_ENERGY_STATUS</b>                        |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_2DH .....   | See Table 2-22  |
| 06_3EH, 06_3FH .....                                 | See Table 2-25  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-28  |
| 06_3F .....  | See Table 2-31  |
| 06_56H, 06_4FH .....                                 | See Table 2-35  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_DRAM_PERF_STATUS</b>                          |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_2DH .....   | See Table 2-22  |
| 06_3EH, 06_3FH .....                                 | See Table 2-25  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-28  |
| 06_3F .....  | See Table 2-31  |
| 06_56H, 06_4FH .....                                 | See Table 2-35  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_DRAM_POWER_INFO</b>                           |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_2DH .....   | See Table 2-22  |
| 06_3EH, 06_3FH .....                                 | See Table 2-25  |
| 06_3F .....  | See Table 2-31  |
| 06_56H, 06_4FH .....                                 | See Table 2-35  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_DRAM_POWER_LIMIT</b>                          |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_2DH .....   | See Table 2-22  |
| 06_3EH, 06_3FH .....                                 | See Table 2-25  |
| 06_3F .....  | See Table 2-31  |
| 06_56H, 06_4FH .....                                 | See Table 2-35  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_EBC_FREQUENCY_ID</b>                          |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_EBC_HARD_POWERON</b>                          |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel                | Location       |
|--|----------------|
| 0FH .....  | See Table 2-45 |
| MSR_EBC_SOFT_POWERON   |                |
| 0FH .....  | See Table 2-45 |
| MSR_EBL_CR_POWERON   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....                 | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH ..... | See Table 2-6  |
| 06_0EH .....   | See Table 2-48 |
| 06_09H .....   | See Table 2-49 |
| MSR_EFSB_DRDY0   |                |
| 0F_03H, 0F_04H .....   | See Table 2-46 |
| MSR_EFSB_DRDY1   |                |
| 0F_03H, 0F_04H .....   | See Table 2-46 |
| MSR_EMON_L3_CTR_CTL0   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL1   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL2   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL3   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL4   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL5   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL6   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_CTR_CTL7   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| 0F_06H .....   | See Table 2-47 |
| MSR_EMON_L3_GL_CTL   |                |
| 06_0FH, 06_17H .....   | See Table 2-3  |
| MSR_ERROR_CONTROL  |                |
| 06_2DH .....   | See Table 2-22 |
| 06_3EH .....   | See Table 2-25 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>         | <b>Location</b> |
|--|-----------------|
| 06_3F .....  | See Table 2-31  |
| <b>MSR_FEATURE_CONFIG</b>                                    |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH ..... | See Table 2-6   |
| 06_25H, 06_2CH .....   | See Table 2-17  |
| 06_2FH .....   | See Table 2-18  |
| 06_2AH, 06_2DH .....   | See Table 2-19  |
| 06_57H .....   | See Table 2-43  |
| <b>MSR_FIRM_ESCRO</b>  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FIRM_ESCR1</b>  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_CCCRO</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_CCCR1</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_CCCR2</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_CCCR3</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_COUNTER0</b>                                    |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_COUNTER1</b>                                    |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_COUNTER2</b>                                    |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_COUNTER3</b>                                    |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_ESCRO</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FLAME_ESCR1</b>                                       |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FSB_ESCRO</b>   |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FSB_ESCR1</b>   |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_FSB_FREQ</b>  |                 |
| 06_0FH, 06_17H .....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....                 | See Table 2-7   |
| 06_4CH .....   | See Table 2-11  |
| 06_0EH .....   | See Table 2-48  |
| <b>MSR_GQ_SNOOP_MESF</b>                                     |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....         | See Table 2-15  |
| MSR_GRAPHICS_PERF_LIMIT_REASONS                      |                 |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| MSR_IFSB_BUSQ0                                       |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IFSB_BUSQ1                                       |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IFSB_CNTR7                                       |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IFSB_CTL6  |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IFSB_SNPQ0                                       |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IFSB_SNPQ1                                       |                 |
| 0F_03H, 0F_04H .....                                 | See Table 2-46  |
| MSR_IQ_CCCRO   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_CCCR1   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_CCCR2   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_CCCR3   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_CCCR4   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_CCCR5   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER0                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER1                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER2                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER3                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER4                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_COUNTER5                                      |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_ESCR0   |                 |
| 0FH .....  | See Table 2-45  |
| MSR_IQ_ESCR1   |                 |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 0FH .....  | See Table 2-45  |
| <b>MSR_IS_ESCR0</b>                                  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_IS_ESCR1</b>                                  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_ITLB_ESCR0</b>                                |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_ITLB_ESCR1</b>                                |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_IX_ESCR0</b>                                  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_IX_ESCR1</b>                                  |                 |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_0</b>                              |                 |
| 0FH .....  | See Table 2-45  |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_0_FROM_IP</b>                      |                 |
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH .....   | See Table 2-12  |
| 06_7AH .....   | See Table 2-13  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_0_TO_IP</b>                        |                 |
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH .....   | See Table 2-12  |
| 06_7AH .....   | See Table 2-13  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_1_FROM_IP</b>                      |                 |
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_1_TO_IP</b>                 |                |
| 06_0FH, 06_17H .....                          | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....  | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....  | See Table 2-7  |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_10_FROM_IP</b>              |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_10_TO_IP</b>                |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_11_FROM_IP</b>              |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_11_TO_IP</b>                |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_12_FROM_IP</b>              |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_12_TO_IP</b>                |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |
| 06_2AH, 06_2DH .....                          | See Table 2-19 |
| 0FH .....                                     | See Table 2-45 |
| <b>MSR_LASTBRANCH_13_FROM_IP</b>              |                |
| 06_5CH, 06_7AH .....                          | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....          | See Table 2-14 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_13_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_14_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_14_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_15_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_15_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_16_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_16_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_17_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_17_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_18_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| <b>MSR_LASTBRANCH_18_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_19_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_19_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_2</b>                              |                 |
| 0FH .....  | See Table 2-45  |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_2_FROM_IP</b>                      |                 |
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_2_TO_IP</b>                        |                 |
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_20_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_20_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_21_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_21_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_22_FROM_IP</b>                     |                 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_22_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_23_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_23_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_24_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_24_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_25_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_25_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_26_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_26_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_27_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_27_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_28_FROM_IP</b>                     |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_28_TO_IP</b>                       |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| <b>MSR_LASTBRANCH_29_FROM_IP</b>                     |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel  | Location       |
|--|----------------|
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_29_TO_IP                        |                |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_3                               |                |
| 0FH  | See Table 2-45 |
| 06_0EH   | See Table 2-48 |
| 06_09H   | See Table 2-49 |
| MSR_LASTBRANCH_3_FROM_IP                       |                |
| 06_0FH, 06_17H                                 | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H         | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH         | See Table 2-7  |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH                 | See Table 2-14 |
| 06_2AH, 06_2DH                                 | See Table 2-19 |
| 0FH  | See Table 2-45 |
| MSR_LASTBRANCH_3_TO_IP                         |                |
| 06_0FH, 06_17H                                 | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H         | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH         | See Table 2-7  |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH                 | See Table 2-14 |
| 06_2AH, 06_2DH                                 | See Table 2-19 |
| 0FH  | See Table 2-45 |
| MSR_LASTBRANCH_30_FROM_IP                      |                |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_30_TO_IP                        |                |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_31_FROM_IP                      |                |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_31_TO_IP                        |                |
| 06_5CH, 06_7AH                                 | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H | See Table 2-38 |
| MSR_LASTBRANCH_4                               |                |
| 06_0EH   | See Table 2-48 |
| 06_09H   | See Table 2-49 |
| MSR_LASTBRANCH_4_FROM_IP                       |                |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H         | See Table 2-4  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_4_TO_IP</b>                        |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_5</b>                              |                 |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_5_FROM_IP</b>                      |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_5_TO_IP</b>                        |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_6</b>                              |                 |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_6_FROM_IP</b>                      |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_6_TO_IP</b>                        |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_7</b>                              |                 |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_7_FROM_IP</b>                      |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_7_TO_IP</b>                        |                 |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_8_FROM_IP</b>                      |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_8_TO_IP</b>                        |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_9_FROM_IP</b>                      |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_9_TO_IP</b>                        |                 |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 0FH .....  | See Table 2-45  |
| <b>MSR_LASTBRANCH_TOS</b>                            |                 |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_0FH, 06_17H .....                                 | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H .....         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH .....         | See Table 2-7   |
| 06_5CH, 06_7AH .....                                 | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_57H .....   | See Table 2-43  |
| 06_0EH .....   | See Table 2-48  |
| 06_09H .....   | See Table 2-49  |
| <b>MSR_LASTBRANCH_INFO_0</b>                         |                 |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_1</b>                                |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_10</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_11</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_12</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_13</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_14</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_15</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_16</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_17</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |
| <b>MSR_LBR_INFO_18</b>                               |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| 06_7AH .....   | See Table 2-13  |

| MSR Name and CPUID DisplayFamily_DisplayModel       | Location       |
|---|----------------|
| MSR_LBR_INFO_19                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_2                                      |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_20                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_21                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_22                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_23                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_24                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_25                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_26                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_27                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_28                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_29                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_3                                      |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |
| MSR_LBR_INFO_30                                     |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H..... | See Table 2-38 |
| 06_7AH.....   | See Table 2-13 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>        | <b>Location</b> |
|---|-----------------|
| <b>MSR_LBR_INFO_31</b>                                      |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_4</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_5</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_6</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_7</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_8</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_INFO_9</b>                                       |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38  |
| 06_7AH.....   | See Table 2-13  |
| <b>MSR_LBR_SELECT</b>                                       |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....                 | See Table 2-7   |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-28  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_LER_FROM_LIP</b>                                     |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| 06_57H.....   | See Table 2-43  |
| 0FH.....  | See Table 2-45  |
| 06_0EH.....   | See Table 2-48  |
| 06_09H.....   | See Table 2-49  |
| <b>MSR_LER_TO_LIP</b>                                       |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_1AH, 06_1EH, 06_1FH, 06_2EH .....                 | See Table 2-14  |
| 06_2AH, 06_2DH .....                                 | See Table 2-19  |
| 06_57H.....  | See Table 2-43  |
| 0FH.....   | See Table 2-45  |
| 06_0EH.....  | See Table 2-48  |
| 06_09H.....  | See Table 2-49  |
| <b>MSR_MO_PMON_ADDR_MASK</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_ADDR_MATCH</b>                        |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_BOX_CTRL</b>                          |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_BOX_OVF_CTRL</b>                      |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_BOX_STATUS</b>                        |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL0</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL1</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL2</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL3</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL4</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_CTRL5</b>                             |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_DSP</b>                               |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL0</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL1</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL2</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL3</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL4</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |
| <b>MSR_MO_PMON_EVTN_SEL5</b>                         |                 |
| 06_2EH.....  | See Table 2-16  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_M0_PMON_ISS<br>06_2EH.....                       | See Table 2-16  |
| MSR_M0_PMON_MAP<br>06_2EH.....                       | See Table 2-16  |
| MSR_M0_PMON_MM_CONFIG<br>06_2EH.....                 | See Table 2-16  |
| MSR_M0_PMON_MSC_THR<br>06_2EH.....                   | See Table 2-16  |
| MSR_M0_PMON_PGT<br>06_2EH.....                       | See Table 2-16  |
| MSR_M0_PMON_PLD<br>06_2EH.....                       | See Table 2-16  |
| MSR_M0_PMON_TIMESTAMP<br>06_2EH.....                 | See Table 2-16  |
| MSR_M0_PMON_ZDP<br>06_2EH.....                       | See Table 2-16  |
| MSR_M1_PMON_ADDR_MASK<br>06_2EH.....                 | See Table 2-16  |
| MSR_M1_PMON_ADDR_MATCH<br>06_2EH.....                | See Table 2-16  |
| MSR_M1_PMON_BOX_CTRL<br>06_2EH.....                  | See Table 2-16  |
| MSR_M1_PMON_BOX_OVF_CTRL<br>06_2EH.....              | See Table 2-16  |
| MSR_M1_PMON_BOX_STATUS<br>06_2EH.....                | See Table 2-16  |
| MSR_M1_PMON_CTR0<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_CTR1<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_CTR2<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_CTR3<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_CTR4<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_CTR5<br>06_2EH.....                      | See Table 2-16  |
| MSR_M1_PMON_DSP<br>06_2EH.....                       | See Table 2-16  |
| MSR_M1_PMON_EVNT_SELO<br>06_2EH.....                 | See Table 2-16  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel                       | Location       |
|---|----------------|
| MSR_M1_PMON_EVT_SEL1<br>06_2EH.....                                 | See Table 2-16 |
| MSR_M1_PMON_EVT_SEL2<br>06_2EH.....                                 | See Table 2-16 |
| MSR_M1_PMON_EVT_SEL3<br>06_2EH.....                                 | See Table 2-16 |
| MSR_M1_PMON_EVT_SEL4<br>06_2EH.....                                 | See Table 2-16 |
| MSR_M1_PMON_EVT_SEL5<br>06_2EH.....                                 | See Table 2-16 |
| MSR_M1_PMON_ISS<br>06_2EH.....                                      | See Table 2-16 |
| MSR_M1_PMON_MAP<br>06_2EH.....                                      | See Table 2-16 |
| MSR_M1_PMON_MM_CONFIG<br>06_2EH.....                                | See Table 2-16 |
| MSR_M1_PMON_MSC_THR<br>06_2EH.....                                  | See Table 2-16 |
| MSR_M1_PMON_PGT<br>06_2EH.....                                      | See Table 2-16 |
| MSR_M1_PMON_PLD<br>06_2EH.....                                      | See Table 2-16 |
| MSR_M1_PMON_TIMESTAMP<br>06_2EH.....                                | See Table 2-16 |
| MSR_M1_PMON_ZDP<br>06_2EH.....                                      | See Table 2-16 |
| IA32_MCO_MISC / MSR_MCO_MISC<br>06_1AH, 06_1EH, 06_1FH, 06_2EH..... | See Table 2-14 |
| MSR_MCO_RESIDENCY<br>06_57H.....                                    | See Table 2-43 |
| IA32_MC1_MISC / MSR_MC1_MISC<br>06_1AH, 06_1EH, 06_1FH, 06_2EH..... | See Table 2-14 |
| IA32_MC10_ADDR / MSR_MC10_ADDR<br>06_2EH.....                       | See Table 2-16 |
| 06_2DH.....   | See Table 2-22 |
| 06_3EH.....   | See Table 2-25 |
| 06_3F.....  | See Table 2-31 |
| 06_56H, 06_4FH.....   | See Table 2-36 |
| 06_4FH.....   | See Table 2-37 |
| IA32_MC10_CTL / MSR_MC10_CTL<br>06_2EH.....                         | See Table 2-16 |
| 06_2DH.....   | See Table 2-22 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC10_MISC / MSR_MC10_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC10_STATUS / MSR_MC10_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC11_ADDR / MSR_MC11_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC11_CTL / MSR_MC11_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC11_MISC / MSR_MC11_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC11_STATUS / MSR_MC11_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| <b>IA32_MC12_ADDR / MSR_MC12_ADDR</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC12_CTL / MSR_MC12_CTL</b>           |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC12_MISC / MSR_MC12_MISC</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC12_STATUS / MSR_MC12_STATUS</b>     |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC13_ADDR / MSR_MC13_ADDR</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC13_CTL / MSR_MC13_CTL</b>           |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC13_MISC / MSR_MC13_MISC</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| <b>IA32_MC13_STATUS / MSR_MC13_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC14_ADDR / MSR_MC14_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC14_CTL / MSR_MC14_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC14_MISC / MSR_MC14_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC14_STATUS / MSR_MC14_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC15_ADDR / MSR_MC15_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC15_CTL / MSR_MC15_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| <b>IA32_MC15_MISC / MSR_MC15_MISC</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC15_STATUS / MSR_MC15_STATUS</b>     |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC16_ADDR / MSR_MC16_ADDR</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC16_CTL / MSR_MC16_CTL</b>           |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC16_MISC / MSR_MC16_MISC</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC16_STATUS / MSR_MC16_STATUS</b>     |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_4FH.....                                   | See Table 2-37 |
| <b>IA32_MC17_ADDR / MSR_MC17_ADDR</b>         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC17_CTL / MSR_MC17_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC17_MISC / MSR_MC17_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC17_STATUS / MSR_MC17_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC18_ADDR / MSR_MC18_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC18_CTL / MSR_MC18_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC18_MISC / MSR_MC18_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC18_STATUS / MSR_MC18_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC19_ADDR / MSR_MC19_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC19_CTL / MSR_MC19_CTL</b>                  |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC19_MISC / MSR_MC19_MISC</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC19_STATUS / MSR_MC19_STATUS</b>            |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC2_MISC / MSR_MC2_MISC</b>                  |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| <b>IA32_MC20_ADDR / MSR_MC20_ADDR</b>                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_4FH.....  | See Table 2-37  |
| IA32_MC20_CTL / MSR_MC20_CTL                         |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| IA32_MC20_MISC / MSR_MC20_MISC                       |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| IA32_MC20_STATUS / MSR_MC20_STATUS                   |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| IA32_MC21_ADDR / MSR_MC21_ADDR                       |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4F.....   | See Table 2-37  |
| IA32_MC21_CTL / MSR_MC21_CTL                         |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4F.....   | See Table 2-37  |
| IA32_MC21_MISC / MSR_MC21_MISC                       |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4F.....   | See Table 2-37  |
| IA32_MC21_STATUS / MSR_MC21_STATUS                   |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4F.....   | See Table 2-37  |
| IA32_MC22_ADDR / MSR_MC22_ADDR                       |                 |
| 06_3EH.....  | See Table 2-25  |
| IA32_MC22_CTL / MSR_MC22_CTL                         |                 |
| 06_3EH.....  | See Table 2-25  |
| IA32_MC22_MISC / MSR_MC22_MISC                       |                 |
| 06_3EH.....  | See Table 2-25  |

MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel     | Location       |
|---|----------------|
| IA32_MC22_STATUS / MSR_MC22_STATUS<br>06_3EH..... | See Table 2-25 |
| IA32_MC23_ADDR / MSR_MC23_ADDR<br>06_3EH.....     | See Table 2-25 |
| IA32_MC23_CTL / MSR_MC23_CTL<br>06_3EH.....       | See Table 2-25 |
| IA32_MC23_MISC / MSR_MC23_MISC<br>06_3EH.....     | See Table 2-25 |
| IA32_MC23_STATUS / MSR_MC23_STATUS<br>06_3EH..... | See Table 2-25 |
| IA32_MC24_ADDR / MSR_MC24_ADDR<br>06_3EH.....     | See Table 2-25 |
| IA32_MC24_CTL / MSR_MC24_CTL<br>06_3EH.....       | See Table 2-25 |
| IA32_MC24_MISC / MSR_MC24_MISC<br>06_3EH.....     | See Table 2-25 |
| IA32_MC24_STATUS / MSR_MC24_STATUS<br>06_3EH..... | See Table 2-25 |
| IA32_MC25_ADDR / MSR_MC25_ADDR<br>06_3EH.....     | See Table 2-25 |
| IA32_MC25_CTL / MSR_MC25_CTL<br>06_3EH.....       | See Table 2-25 |
| IA32_MC25_MISC / MSR_MC25_MISC<br>06_3EH.....     | See Table 2-25 |
| IA32_MC25_STATUS / MSR_MC25_STATUS<br>06_3EH..... | See Table 2-25 |
| IA32_MC26_ADDR / MSR_MC26_ADDR<br>06_3EH.....     | See Table 2-25 |
| IA32_MC26_CTL / MSR_MC26_CTL<br>06_3EH.....       | See Table 2-25 |
| IA32_MC26_MISC / MSR_MC26_MISC<br>06_3EH.....     | See Table 2-25 |
| IA32_MC26_STATUS / MSR_MC26_STATUS<br>06_3EH..... | See Table 2-25 |
| IA32_MC27_ADDR / MSR_MC27_ADDR<br>06_3EH.....     | See Table 2-25 |
| IA32_MC27_CTL / MSR_MC27_CTL<br>06_3EH.....       | See Table 2-25 |
| IA32_MC27_MISC / MSR_MC27_MISC<br>06_3EH.....     | See Table 2-25 |
| IA32_MC27_STATUS / MSR_MC27_STATUS<br>06_3EH..... | See Table 2-25 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>        | <b>Location</b> |
|---|-----------------|
| IA32_MC28_ADDR / MSR_MC28_ADDR                              |                 |
| 06_3EH.....   | See Table 2-25  |
| IA32_MC28_CTL / MSR_MC28_CTL                                |                 |
| 06_3EH.....   | See Table 2-25  |
| IA32_MC28_MISC / MSR_MC28_MISC                              |                 |
| 06_3EH.....   | See Table 2-25  |
| IA32_MC28_STATUS / MSR_MC28_STATUS                          |                 |
| 06_3EH.....   | See Table 2-25  |
| IA32_MC29_ADDR / MSR_MC29_ADDR                              |                 |
| 06_3EH.....   | See Table 2-26  |
| IA32_MC29_CTL / MSR_MC29_CTL                                |                 |
| 06_3EH.....   | See Table 2-26  |
| IA32_MC29_MISC / MSR_MC29_MISC                              |                 |
| 06_3EH.....   | See Table 2-26  |
| IA32_MC29_STATUS / MSR_MC29_STATUS                          |                 |
| 06_3EH.....   | See Table 2-26  |
| IA32_MC3_ADDR / MSR_MC3_ADDR                                |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_57H.....   | See Table 2-43  |
| 06_0EH.....   | See Table 2-48  |
| 06_09H.....   | See Table 2-49  |
| IA32_MC3_CTL / MSR_MC3_CTL                                  |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_57H.....   | See Table 2-43  |
| 06_0EH.....   | See Table 2-48  |
| 06_09H.....   | See Table 2-49  |
| IA32_MC3_MISC / MSR_MC3_MISC                                |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_0EH.....   | See Table 2-48  |
| IA32_MC3_STATUS / MSR_MC3_STATUS                            |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_57H.....   | See Table 2-43  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel               | Location       |
|---|----------------|
| 06_0EH.....   | See Table 2-48 |
| 06_09H.....   | See Table 2-49 |
| IA32_MC30_ADDR / MSR_MC30_ADDR                              |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC30_CTL / MSR_MC30_CTL                                |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC30_MISC / MSR_MC30_MISC                              |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC30_STATUS / MSR_MC30_STATUS                          |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC31_ADDR / MSR_MC31_ADDR                              |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC31_CTL / MSR_MC31_CTL                                |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC31_MISC / MSR_MC31_MISC                              |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC31_STATUS / MSR_MC31_STATUS                          |                |
| 06_3EH.....   | See Table 2-26 |
| IA32_MC4_ADDR / MSR_MC4_ADDR                                |                |
| 06_0FH, 06_17H.....   | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14 |
| 06_57H.....   | See Table 2-43 |
| 06_0EH.....   | See Table 2-48 |
| 06_09H.....   | See Table 2-49 |
| IA32_MC4_CTL / MSR_MC4_CTL                                  |                |
| 06_0FH, 06_17H.....   | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14 |
| 06_57H.....   | See Table 2-43 |
| 06_0EH.....   | See Table 2-48 |
| 06_09H.....   | See Table 2-49 |
| IA32_MC4_CTL2 / MSR_MC4_CTL2                                |                |
| 06_2AH, 06_2DH.....   | See Table 2-19 |
| IA32_MC4_STATUS / MSR_MC4_STATUS                            |                |
| 06_0FH, 06_17H.....   | See Table 2-3  |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4  |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14 |
| 06_57H.....   | See Table 2-43 |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>        | <b>Location</b> |
|---|-----------------|
| 06_0EH.....   | See Table 2-48  |
| 06_09H.....   | See Table 2-49  |
| <b>MSR_MC5_ADDR / MSR_MC5_ADDR</b>                          |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2DH.....   | See Table 2-22  |
| 06_3EH.....   | See Table 2-25  |
| 06_3FH.....   | See Table 2-31  |
| 06_4FH.....   | See Table 2-37  |
| 06_57H.....   | See Table 2-43  |
| 06_0EH.....   | See Table 2-48  |
| <b>IA32_MC5_CTL / MSR_MC5_CTL</b>                           |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2DH.....   | See Table 2-22  |
| 06_3EH.....   | See Table 2-25  |
| 06_3FH.....   | See Table 2-31  |
| 06_4FH.....   | See Table 2-37  |
| 06_57H.....   | See Table 2-43  |
| 06_0EH.....   | See Table 2-48  |
| <b>IA32_MC5_MISC / MSR_MC5_MISC</b>                         |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2DH.....   | See Table 2-22  |
| 06_3EH.....   | See Table 2-25  |
| 06_3FH.....   | See Table 2-31  |
| 06_4FH.....   | See Table 2-37  |
| 06_0EH.....   | See Table 2-48  |
| <b>IA32_MC5_STATUS / MSR_MC5_STATUS</b>                     |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2DH.....   | See Table 2-22  |
| 06_3EH.....   | See Table 2-25  |
| 06_3FH.....   | See Table 2-31  |
| 06_4FH.....   | See Table 2-37  |
| 06_57H.....   | See Table 2-43  |
| 06_0EH.....   | See Table 2-48  |
| <b>IA32_MC6_ADDR / MSR_MC6_ADDR</b>                         |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC6_CTL / MSR_MC6_CTL                    |                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| MSR_MC6_DEMOTION_POLICY_CONFIG                |                |
| 06_37H.....                                   | See Table 2-9  |
| IA32_MC6_MISC / MSR_MC6_MISC                  |                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| MSR_MC6_RESIDENCY_COUNTER                     |                |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....   | See Table 2-7  |
| 06_37H.....                                   | See Table 2-9  |
| 06_57H.....                                   | See Table 2-43 |
| IA32_MC6_STATUS / MSR_MC6_STATUS              |                |
| 06_0FH, 06_17H.....                           | See Table 2-3  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3FH.....                                   | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC7_ADDR / MSR_MC7_ADDR                  |                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC7_CTL / MSR_MC7_CTL                    |                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC7_MISC / MSR_MC7_MISC</b>                  |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC7_STATUS / MSR_MC7_STATUS</b>              |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_56H, 06_4FH.....                                  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC8_ADDR / MSR_MC8_ADDR</b>                  |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC8_CTL / MSR_MC8_CTL</b>                    |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC8_MISC / MSR_MC8_MISC</b>                  |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |
| 06_4FH.....  | See Table 2-37  |
| <b>IA32_MC8_STATUS / MSR_MC8_STATUS</b>              |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2DH.....  | See Table 2-22  |
| 06_3EH.....  | See Table 2-25  |
| 06_3F.....   | See Table 2-31  |

MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC9_ADDR / MSR_MC9_ADDR                  |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC9_CTL / MSR_MC9_CTL                    |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC9_MISC / MSR_MC9_MISC                  |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| IA32_MC9_STATUS / MSR_MC9_STATUS              |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_2DH.....                                   | See Table 2-22 |
| 06_3EH.....                                   | See Table 2-25 |
| 06_3F.....                                    | See Table 2-31 |
| 06_56H, 06_4FH.....                           | See Table 2-36 |
| 06_4FH.....                                   | See Table 2-37 |
| MSR_MCG_MISC                                  |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R10                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R11                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R12                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R13                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R14                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R15                                   |                |

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R8                                    |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_R9                                    |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RAX                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RBP                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RBX                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RCX                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RDI                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RDX                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RESERVED1 - MSR_MCG_RESERVED5         |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RFLAGS                                |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RIP                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RSI                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MCG_RSP                                   |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MISC_FEATURE_CONTROL                      |                |
| 06_5CH, 06_7AH.....                           | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2AH, 06_2DH.....                           | See Table 2-19 |
| MSR_MISC_PWR_MGMT                             |                |
| 06_5CH, 06_7AH.....                           | See Table 2-12 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....           | See Table 2-14 |
| 06_2AH, 06_2DH.....                           | See Table 2-19 |
| MSR_MOB_ESCRO                                 |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MOB_ESCR1                                 |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MS_CCCRO                                  |                |
| 0FH.....                                      | See Table 2-45 |
| MSR_MS_CCCR1                                  |                |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel               | Location       |
|---|----------------|
| 0FH.....  | See Table 2-45 |
| MSR_MS_CCCR2  |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_CCCR3  |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_COUNTER0   |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_COUNTER1   |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_COUNTER2   |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_COUNTER3   |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_ESCRO  |                |
| 0FH.....  | See Table 2-45 |
| MSR_MS_ESCR1  |                |
| 0FH.....  | See Table 2-45 |
| MSR_MTRRCAP   |                |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....         | See Table 2-38 |
| MSR_OFFCORE_RSP_0   |                |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14 |
| 06_2AH, 06_2DH.....   | See Table 2-19 |
| 06_57H.....   | See Table 2-43 |
| MSR_OFFCORE_RSP_1   |                |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_25H, 06_2CH.....   | See Table 2-17 |
| 06_2FH.....   | See Table 2-18 |
| 06_2AH, 06_2DH.....   | See Table 2-19 |
| 06_57H.....   | See Table 2-43 |
| MSR_PCIE_PLL_RATIO  |                |
| 06_3FH.....   | See Table 2-31 |
| MSR_PCU_PMON_BOX_CTL  |                |
| 06_2DH.....   | See Table 2-23 |
| 06_3FH.....   | See Table 2-32 |
| MSR_PCU_PMON_BOX_FILTER                                     |                |
| 06_2DH.....   | See Table 2-23 |
| 06_3FH.....   | See Table 2-32 |
| MSR_PCU_PMON_BOX_STATUS                                     |                |
| 06_3EH.....   | See Table 2-27 |
| 06_3FH.....   | See Table 2-32 |
| MSR_PCU_PMON_CTRL0  |                |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_CTR1</b>                             |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_CTR2</b>                             |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_CTR3</b>                             |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_EVNTSELO</b>                         |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_EVNTSEL1</b>                         |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_EVNTSEL2</b>                         |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PCU_PMON_EVNTSEL3</b>                         |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| <b>MSR_PEBS_ENABLE</b>                               |                 |
| 06_0FH, 06_17H.....                                  | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....          | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....          | See Table 2-7   |
| 06_5CH.....  | See Table 2-12  |
| 06_7AH.....  | See Table 2-13  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2AH, 06_2DH.....                                  | See Table 2-19  |
| 06_3EH.....  | See Table 2-26  |
| 06_57H.....  | See Table 2-43  |
| 0FH.....   | See Table 2-45  |
| <b>MSR_PEBS_FRONTEND</b>                             |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....  | See Table 2-38  |
| <b>MSR_PEBS_LD_LAT</b>                               |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                  | See Table 2-14  |
| 06_2AH, 06_2DH.....                                  | See Table 2-19  |
| <b>MSR_PEBS_MATRIX_VERT</b>                          |                 |
| 0FH.....   | See Table 2-45  |
| <b>MSR_PEBS_NUM_ALT</b>                              |                 |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel                       | Location                         |
|---|----------------------------------|
| 06_2DH.....   | See Table 2-22                   |
| MSR_PERF_CAPABILITIES   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_FIXED_CTR_CTRL   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_FIXED_CTR0   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_FIXED_CTR1   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_FIXED_CTR2   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_GLOBAL_CTRL  |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| MSR_PERF_GLOBAL_OVF_CTRL  |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14                   |
| MSR_PERF_GLOBAL_STATUS  |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14                   |
| MSR_PERF_STATUS   |                                  |
| 06_0FH, 06_17H.....   | See Table 2-3                    |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                         | See Table 2-4                    |
| 06_2AH, 06_2DH.....   | See Table 2-19                   |
| MSR_PKG_C10_RESIDENCY   |                                  |
| 06_5CH, 06_7AH.....   | See Table 2-12                   |
| 06_45H.....   | See Table 2-29 and<br>Table 2-30 |
| 06_4FH.....   | See Table 2-37                   |
| MSR_PKG_C2_RESIDENCY  |                                  |
| 06_27H.....   | See Table 2-5                    |
| 06_5CH, 06_7AH.....   | See Table 2-12                   |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19                   |
| 06_57H.....   | See Table 2-43                   |
| MSR_PKG_C3_RESIDENCY  |                                  |
| 06_5CH, 06_7AH.....   | See Table 2-12                   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH, 06_2FH.....         | See Table 2-14                   |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19                   |
| 06_66H.....   | See Table 2-41                   |
| 06_57H.....   | See Table 2-43                   |
| MSR_PKG_C4_RESIDENCY  |                                  |
| 06_27H.....   | See Table 2-5                    |
| MSR_PKG_C6_RESIDENCY  |                                  |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                | <b>Location</b> |
|---|-----------------|
| 06_27H.....   | See Table 2-5   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....                         | See Table 2-7   |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH, 06_2FH.....         | See Table 2-14  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKG_C7_RESIDENCY</b>   |                 |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH, 06_2FH.....         | See Table 2-14  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKG_C8_RESIDENCY</b>   |                 |
| 06_45H.....   | See Table 2-30  |
| 06_4FH.....   | See Table 2-37  |
| <b>MSR_PKG_C9_RESIDENCY</b>   |                 |
| 06_45H.....   | See Table 2-30  |
| 06_4FH.....   | See Table 2-37  |
| <b>MSR_PKG_CST_CONFIG_CONTROL</b>                                   |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....                         | See Table 2-7   |
| 06_4CH.....   | See Table 2-11  |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| 06_3AH.....   | See Table 2-24  |
| 06_3EH.....   | See Table 2-25  |
| 06_3CH, 06_45H, 06_46H.....   | See Table 2-29  |
| 06_45H.....   | See Table 2-30  |
| 06_3F.....  | See Table 2-31  |
| 06_3DH.....   | See Table 2-34  |
| 06_56H, 06_4FH.....   | See Table 2-35  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKG_ENERGY_STATUS</b>  |                 |
| 06_37H, 06_4AH, 06_5AH, 06_5DH.....                                 | See Table 2-8   |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| <b>MSR_PKG_HDC_CONFIG</b>   |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                 | See Table 2-38  |
| <b>MSR_PKG_HDC_DEEP_RESIDENCY</b>                                   |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                 | See Table 2-38  |
| <b>MSR_PKG_HDC_SHALLOW_RESIDENCY</b>                                |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                 | See Table 2-38  |
| <b>MSR_PKG_PERF_STATUS</b>  |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                | <b>Location</b> |
|---|-----------------|
| 06_2DH.....   | See Table 2-22  |
| 06_3EH, 06_3FH.....   | See Table 2-25  |
| 06_3CH, 06_45H, 06_46H.....   | See Table 2-29  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKG_POWER_INFO</b>   |                 |
| 06_4DH.....   | See Table 2-10  |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKG_POWER_LIMIT</b>  |                 |
| 06_37H, 06_4AH, 06_5AH, 06_5DH.....                                 | See Table 2-8   |
| 06_4DH.....   | See Table 2-10  |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_PKGC_IRTL1</b>   |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....   | See Table 2-28  |
| <b>MSR_PKGC_IRTL2</b>   |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....   | See Table 2-28  |
| <b>MSR_PKGC3_IRTL</b>   |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| <b>MSR_PKGC6_IRTL</b>   |                 |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| <b>MSR_PKGC7_IRTL</b>   |                 |
| 06_2AH.....   | See Table 2-20  |
| <b>MSR_PLATFORM_BRV</b>   |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_PLATFORM_ENERGY_COUNTER</b>                                  |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                 | See Table 2-38  |
| <b>MSR_PLATFORM_ID</b>  |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                         | See Table 2-4   |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH.....                         | See Table 2-7   |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14  |
| <b>MSR_PLATFORM_INFO</b>  |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                | <b>Location</b>               |
|---|-------------------------------|
| 06_3AH.....   | See Table 2-24                |
| 06_3EH.....   | See Table 2-25                |
| 06_3CH, 06_45H, 06_46H.....   | See Table 2-28 and Table 2-29 |
| 06_56H, 06_4FH.....   | See Table 2-35                |
| 06_57H.....   | See Table 2-43                |
| <b>MSR_PLATFORM_POWER_LIMIT</b>                                     |                               |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                 | See Table 2-38                |
| <b>MSR_PMG_IO_CAPTURE_BASE</b>                                      |                               |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH.....         | See Table 2-6                 |
| 06_4CH.....   | See Table 2-11                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14                |
| 06_2AH, 06_2DH.....   | See Table 2-19                |
| 06_3AH.....   | See Table 2-24                |
| 06_3EH.....   | See Table 2-25                |
| 06_57H.....   | See Table 2-43                |
| <b>MSR_PMH_ESCRO</b>  |                               |
| 0FH.....  | See Table 2-45                |
| <b>MSR_PMH_ESCR1</b>  |                               |
| 0FH.....  | See Table 2-45                |
| <b>MSR_PMON_GLOBAL_CONFIG</b>                                       |                               |
| 06_3EH.....   | See Table 2-27                |
| 06_3FH.....   | See Table 2-32                |
| <b>MSR_PMON_GLOBAL_CTL</b>  |                               |
| 06_3EH.....   | See Table 2-27                |
| 06_3FH.....   | See Table 2-32                |
| <b>MSR_PMON_GLOBAL_STATUS</b>                                       |                               |
| 06_3EH.....   | See Table 2-27                |
| 06_3FH.....   | See Table 2-32                |
| <b>MSR_POWER_CTL</b>  |                               |
| 06_5CH, 06_7AH.....   | See Table 2-12                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                                 | See Table 2-14                |
| 06_2AH, 06_2DH.....   | See Table 2-19                |
| <b>MSR_PPO_ENERGY_STATUS</b>  |                               |
| 06_37H, 06_4AH, 06_5AH, 06_5DH.....                                 | See Table 2-8                 |
| 06_5CH, 06_7AH.....   | See Table 2-12                |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19                |
| 06_57H.....   | See Table 2-43                |
| <b>MSR_PPO_POLICY</b>   |                               |
| 06_2AH, 06_45H.....   | See Table 2-20                |
| <b>MSR_PPO_POWER_LIMIT</b>  |                               |
| 06_4CH.....   | See Table 2-11                |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>           | <b>Location</b> |
|--|-----------------|
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H | See Table 2-19  |
| 06_57H   | See Table 2-43  |
| <b>MSR_PP1_ENERGY_STATUS</b>                                   |                 |
| 06_5CH, 06_7AH   | See Table 2-12  |
| 06_2AH, 06_45H   | See Table 2-20  |
| 06_3CH, 06_45H, 06_46H   | See Table 2-29  |
| <b>MSR_PP1_POLICY</b>  |                 |
| 06_2AH, 06_45H   | See Table 2-20  |
| 06_3CH, 06_45H, 06_46H   | See Table 2-29  |
| <b>MSR_PP1_POWER_LIMIT</b>                                     |                 |
| 06_2AH, 06_45H   | See Table 2-20  |
| 06_3CH, 06_45H, 06_46H   | See Table 2-29  |
| <b>MSR_PPERF</b>   |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H                 | See Table 2-38  |
| <b>MSR_PPIN</b>  |                 |
| 06_3EH   | See Table 2-25  |
| 06_56H, 06_4FH   | See Table 2-35  |
| <b>MSR_PPIN_CTL</b>  |                 |
| 06_3EH   | See Table 2-25  |
| 06_56H, 06_4FH   | See Table 2-35  |
| <b>MSR_PRMRR_PHYS_BASE</b>                                     |                 |
| 06_8EH, 06_9EH   | See Table 2-40  |
| <b>MSR_PRMRR_PHYS_MASK</b>                                     |                 |
| 06_8EH, 06_9EH   | See Table 2-40  |
| <b>MSR_PRMRR_VALID_CONFIG</b>                                  |                 |
| 06_8EH, 06_9EH   | See Table 2-40  |
| <b>MSR_RING_RATIO_LIMIT</b>                                    |                 |
| 06_8EH, 06_9EH   | See Table 2-40  |
| <b>MSR_RO_PMON_BOX_CTRL</b>                                    |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_BOX_OVF_CTRL</b>                                |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_BOX_STATUS</b>                                  |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_CTRL0</b>                                       |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_CTRL1</b>                                       |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_CTRL2</b>                                       |                 |
| 06_2EH   | See Table 2-16  |
| <b>MSR_RO_PMON_CTRL3</b>                                       |                 |
| 06_2EH   | See Table 2-16  |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_RO_PMON_CTR4<br>06_2EH.....                      | See Table 2-16  |
| MSR_RO_PMON_CTR5<br>06_2EH.....                      | See Table 2-16  |
| MSR_RO_PMON_CTR6<br>06_2EH.....                      | See Table 2-16  |
| MSR_RO_PMON_CTR7<br>06_2EH.....                      | See Table 2-16  |
| MSR_RO_PMON_EVNT_SELO<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL1<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL2<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL3<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL4<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL5<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL6<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_EVNT_SEL7<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P0<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P1<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P2<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P3<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P4<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P5<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P6<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_IPERFO_P7<br>06_2EH.....                 | See Table 2-16  |
| MSR_RO_PMON_QLX_P0<br>06_2EH.....                    | See Table 2-16  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| MSR_R0_PMON_QLX_P1<br>06_2EH.....             | See Table 2-16 |
| MSR_R0_PMON_QLX_P2<br>06_2EH.....             | See Table 2-16 |
| MSR_R0_PMON_QLX_P3<br>06_2EH.....             | See Table 2-16 |
| MSR_R1_PMON_BOX_CTRL<br>06_2EH.....           | See Table 2-16 |
| MSR_R1_PMON_BOX_OVF_CTRL<br>06_2EH.....       | See Table 2-16 |
| MSR_R1_PMON_BOX_STATUS<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_CTR10<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR11<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR12<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR13<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR14<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR15<br>06_2EH.....              | See Table 2-16 |
| MSR_R1_PMON_CTR8<br>06_2EH.....               | See Table 2-16 |
| MSR_R1_PMON_CTR9<br>06_2EH.....               | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL10<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL11<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL12<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL13<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL14<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL15<br>06_2EH.....         | See Table 2-16 |
| MSR_R1_PMON_EVNT_SEL8<br>06_2EH.....          | See Table 2-16 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                | <b>Location</b> |
|---|-----------------|
| MSR_R1_PMON_EVNT_SEL9<br>06_2EH.....                                | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P10<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P11<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P12<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P13<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P14<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P15<br>06_2EH.....                               | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P8<br>06_2EH.....                                | See Table 2-16  |
| MSR_R1_PMON_IPERF1_P9<br>06_2EH.....                                | See Table 2-16  |
| MSR_R1_PMON_QLX_P4<br>06_2EH.....                                   | See Table 2-16  |
| MSR_R1_PMON_QLX_P5<br>06_2EH.....                                   | See Table 2-16  |
| MSR_R1_PMON_QLX_P6<br>06_2EH.....                                   | See Table 2-16  |
| MSR_R1_PMON_QLX_P7<br>06_2EH.....                                   | See Table 2-16  |
| MSR_RAPL_POWER_UNIT<br>06_37H, 06_4AH, 06_5AH, 06_5DH .....         | See Table 2-8   |
| 06_4DH.....   | See Table 2-10  |
| 06_5CH, 06_7AH .....  | See Table 2-12  |
| 06_2AH, 06_2DH, 06_3AH, 06_3CH, 06_3EH, 06_3FH, 06_45H, 06_46H..... | See Table 2-19  |
| 06_3FH.....   | See Table 2-31  |
| 06_56H, 06_4FH .....  | See Table 2-35  |
| 06_57H.....   | See Table 2-43  |
| MSR_RAT_ESCR0<br>0FH.....   | See Table 2-45  |
| MSR_RAT_ESCR1<br>0FH.....   | See Table 2-45  |
| MSR_RING_PERF_LIMIT_REASONS<br>06_3CH, 06_45H, 06_46H .....         | See Table 2-29  |
| MSR_S0_PMON_BOX_CTRL<br>06_2EH.....                                 | See Table 2-16  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_BOX_FILTER                        |                |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_BOX_OVF_CTRL                      |                |
| 06_2EH.....                                   | See Table 2-16 |
| MSR_SO_PMON_BOX_STATUS                        |                |
| 06_2EH.....                                   | See Table 2-16 |
| MSR_SO_PMON_CTR0                              |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_CTR1                              |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_CTR2                              |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_CTR3                              |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_EVNT_SEL0                         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_EVNT_SEL1                         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_EVNT_SEL2                         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_EVNT_SEL3                         |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_SO_PMON_MASK                              |                |
| 06_2EH.....                                   | See Table 2-16 |
| MSR_SO_PMON_MATCH                             |                |
| 06_2EH.....                                   | See Table 2-16 |
| MSR_S1_PMON_BOX_CTRL                          |                |
| 06_2EH.....                                   | See Table 2-16 |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_S1_PMON_BOX_FILTER                        |                |
| 06_3FH.....                                   | See Table 2-32 |
| MSR_S1_PMON_BOX_OVF_CTRL                      |                |
| 06_2EH.....                                   | See Table 2-16 |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_S1_PMON_BOX_STATUS                               |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_S1_PMON_CTRL0                                    |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_CTRL1                                    |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_CTRL2                                    |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_CTRL3                                    |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_EVTN_SEL0                                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_EVTN_SEL1                                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_EVTN_SEL2                                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_EVTN_SEL3                                |                 |
| 06_2EH.....  | See Table 2-16  |
| 06_3FH.....  | See Table 2-32  |
| MSR_S1_PMON_MASK                                     |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_S1_PMON_MATCH                                    |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_S2_PMON_BOX_CTL                                  |                 |
| 06_3FH.....  | See Table 2-32  |
| MSR_S2_PMON_BOX_FILTER                               |                 |
| 06_3FH.....  | See Table 2-32  |
| MSR_S2_PMON_CTRL0                                    |                 |
| 06_3FH.....  | See Table 2-32  |
| MSR_S2_PMON_CTRL1                                    |                 |
| 06_3FH.....  | See Table 2-32  |
| MSR_S2_PMON_CTRL2                                    |                 |
| 06_3FH.....  | See Table 2-32  |
| MSR_S2_PMON_CTRL3                                    |                 |
| 06_3FH.....  | See Table 2-32  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel                                | Location       |
|--|----------------|
| MSR_S2_PMON_EVTNSELO<br>06_3FH.....  | See Table 2-32 |
| MSR_S2_PMON_EVTNSEL1<br>06_3FH.....  | See Table 2-32 |
| MSR_S2_PMON_EVTNSEL2<br>06_3FH.....  | See Table 2-32 |
| MSR_S2_PMON_EVTNSEL3<br>06_3FH.....  | See Table 2-32 |
| MSR_S3_PMON_BOX_CTL<br>06_3FH.....   | See Table 2-32 |
| MSR_S3_PMON_BOX_FILTER<br>06_3FH.....  | See Table 2-32 |
| MSR_S3_PMON_CTRL0<br>06_3FH.....   | See Table 2-32 |
| MSR_S3_PMON_CTRL1<br>06_3FH.....   | See Table 2-32 |
| MSR_S3_PMON_CTRL2<br>06_3FH.....   | See Table 2-32 |
| MSR_S3_PMON_CTRL3<br>06_3FH.....   | See Table 2-32 |
| MSR_S3_PMON_EVTNSELO<br>06_3FH.....  | See Table 2-32 |
| MSR_S3_PMON_EVTNSEL1<br>06_3FH.....  | See Table 2-32 |
| MSR_S3_PMON_EVTNSEL2<br>06_3FH.....  | See Table 2-32 |
| MSR_S3_PMON_EVTNSEL3<br>06_3FH.....  | See Table 2-32 |
| MSR_SAAT_ESCR0<br>0FH.....   | See Table 2-45 |
| MSR_SAAT_ESCR1<br>0FH.....   | See Table 2-45 |
| MSR_SGXOWNEREPOCH0<br>06_5CH, 06_7AH.....                                    | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                          | See Table 2-38 |
| MSR_SGXOWNEREPOCH1<br>06_5CH, 06_7AH.....                                    | See Table 2-12 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H.....                          | See Table 2-38 |
| MSR_SMI_COUNT<br>06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6  |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....  | See Table 2-14 |
| 06_2AH, 06_2DH.....  | See Table 2-19 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>        | <b>Location</b> |
|---|-----------------|
| 06_57H.....   | See Table 2-43  |
| <b>MSR_SMM_BLOCKED</b>                                      |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-29  |
| <b>MSR_SMM_DELAYED</b>                                      |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-29  |
| <b>MSR_SMM_FEATURE_CONTROL</b>                              |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-29  |
| <b>MSR_SMM_MCA_CAP</b>                                      |                 |
| 06_5CH, 06_7AH.....   | See Table 2-12  |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-29  |
| 06_3FH.....   | See Table 2-31  |
| 06_56H, 06_4FH.....   | See Table 2-35  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_SMRR_PHYSBASE</b>                                    |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| <b>MSR_SMRR_PHYSMASK</b>                                    |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| <b>MSR_SSU_ESCR0</b>  |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TBPU_ESCR0</b>                                       |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TBPU_ESCR1</b>                                       |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TC_ESCR0</b>   |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TC_ESCR1</b>   |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TC_PRECISE_EVENT</b>                                 |                 |
| 0FH.....  | See Table 2-45  |
| <b>MSR_TEMPERATURE_TARGET</b>                               |                 |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6   |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14  |
| 06_2AH, 06_2DH.....   | See Table 2-19  |
| 06_3EH.....   | See Table 2-25  |
| 06_56H, 06_4FH.....   | See Table 2-35  |
| 06_57H.....   | See Table 2-43  |
| <b>MSR_THERM2_CTL</b>                                       |                 |
| 06_0FH, 06_17H.....   | See Table 2-3   |
| 06_1CH, 06_26H, 06_27H, 06_35H, 06_36H.....                 | See Table 2-4   |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel               | Location                      |
|---|-------------------------------|
| 0FH.....  | See Table 2-45                |
| 06_0EH.....   | See Table 2-48                |
| 06_09H.....   | See Table 2-49                |
| MSR_THREAD_ID_INFO  |                               |
| 06_3FH.....   | See Table 2-31                |
| MSR_TRACE_HUB_STH ACPIBAR_BASE                              |                               |
| 06_8EH, 06_9EH.....   | See Table 2-40                |
| MSR_TURBO_ACTIVATION_RATIO                                  |                               |
| 06_5CH, 06_7AH.....   | See Table 2-12                |
| 06_3AH.....   | See Table 2-24                |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-28                |
| 06_57H.....   | See Table 2-43                |
| MSR_TURBO_GROUP_CORECNT                                     |                               |
| 06_5CH, 06_7AH.....   | See Table 2-12                |
| MSR_TURBO_POWER_CURRENT_LIMIT                               |                               |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH.....                         | See Table 2-14                |
| MSR_TURBO_RATIO_LIMIT                                       |                               |
| 06_37H, 06_4AH, 06_4DH, 06_5AH, 06_5DH, 06_5CH, 06_7AH..... | See Table 2-6                 |
| 06_4DH.....   | See Table 2-10                |
| 06_5CH, 06_7AH.....   | See Table 2-12                |
| 06_1AH, 06_1EH, 06_1FH, 06_2EH, 06_25H, 06_2CH.....         | See Table 2-14                |
| 06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH.....                 | See Table 2-15                |
| 06_2EH.....   | See Table 2-16                |
| 06_25H, 06_2CH.....   | See Table 2-17                |
| 06_2FH.....   | See Table 2-18                |
| 06_2AH, 06_45H.....   | See Table 2-20                |
| 06_2DH.....   | See Table 2-22                |
| 06_3EH.....   | See Table 2-25 and Table 2-26 |
| 06_3CH, 06_45H, 06_46H.....                                 | See Table 2-29                |
| 06_3FH.....   | See Table 2-31                |
| 06_3DH.....   | See Table 2-34                |
| 06_56H, 06_4FH.....   | See Table 2-35                |
| 06_55H.....   | See Table 2-42                |
| 06_57H.....   | See Table 2-43                |
| MSR_TURBO_RATIO_LIMIT1                                      |                               |
| 06_3EH.....   | See Table 2-25 and Table 2-26 |
| 06_3FH.....   | See Table 2-31                |
| 06_56H, 06_4FH.....   | See Table 2-35                |
| MSR_TURBO_RATIO_LIMIT2                                      |                               |
| 06_3FH.....   | See Table 2-31                |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| MSR_TURBO_RATIO_LIMIT3                               |                 |
| 06_56H.....  | See Table 2-36  |
| 06_4FH.....  | See Table 2-37  |
| MSR_TURBO_RATIO_LIMIT_CORES                          |                 |
| 06_55H.....  | See Table 2-42  |
| MSR_U_PMON_BOX_STATUS                                |                 |
| 06_3EH.....  | See Table 2-27  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_CTR                                       |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_U_PMON_CTR0                                      |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_CTR1                                      |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_EVNT_SEL                                  |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_U_PMON_EVNTSELO                                  |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_EVNTSEL1                                  |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_GLOBAL_CTRL                               |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_U_PMON_GLOBAL_OVF_CTRL                           |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_U_PMON_GLOBAL_STATUS                             |                 |
| 06_2EH.....  | See Table 2-16  |
| MSR_U_PMON_UCLK_FIXED_CTL                            |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U_PMON_UCLK_FIXED_CTR                            |                 |
| 06_2DH.....  | See Table 2-23  |
| 06_3FH.....  | See Table 2-32  |
| MSR_U2L_ESCR0  |                 |
| 0FH.....   | See Table 2-45  |
| MSR_U2L_ESCR1  |                 |
| 0FH.....   | See Table 2-45  |
| MSR_UNC_ARB_PERFCTRO                                 |                 |
| 06_2AH.....  | See Table 2-21  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_ARB_PERFCTR1                          |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_ARB_PERFEVTSELO                       |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_ARB_PERFEVTSEL1                       |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_CBO_0_PERFCTR0                        |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_CBO_0_PERFCTR1                        |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_CBO_0_PERFCTR2                        |                |
| 06_2AH  | See Table 2-21 |
| MSR_UNC_CBO_0_PERFCTR3                        |                |
| 06_2AH  | See Table 2-21 |
| MSR_UNC_CBO_0_PERFEVTSELO                     |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_CBO_0_PERFEVTSEL1                     |                |
| 06_2AH  | See Table 2-21 |
| 06_3CH, 06_45H, 06_46H                        | See Table 2-29 |
| 06_4EH, 06_5EH                                | See Table 2-39 |
| MSR_UNC_CBO_0_PERFEVTSEL2                     |                |
| 06_2AH  | See Table 2-21 |
| MSR_UNC_CBO_0_PERFEVTSEL3                     |                |
| 06_2AH  | See Table 2-21 |
| MSR_UNC_CBO_0_UNIT_STATUS                     |                |
| 06_2AH  | See Table 2-21 |
| MSR_UNC_CBO_1_PERFCTR0                        |                |
| 06_2AH  | See Table 2-21 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_1_PERFCTR1</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_1_PERFCTR2</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_1_PERFCTR3</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_1_PERFEVTSELO</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_1_PERFEVTSEL1</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_1_PERFEVTSEL2</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_1_PERFEVTSEL3</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_1_UNIT_STATUS</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_2_PERFCTR0</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_2_PERFCTR1</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_2_PERFCTR2</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_2_PERFCTR3</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_2_PERFEVTSELO</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_2_PERFEVTSEL1</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_2_PERFEVTSEL2</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_2_PERFEVTSEL3</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_2_UNIT_STATUS</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_3_PERFCTR0</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_3_PERFCTR1</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_3_PERFCTR2</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_3_PERFCTR3</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_3_PERFEVTSELO</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_3_PERFEVTSEL1</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....                         | See Table 2-29  |
| 06_4EH, 06_5EH .....                                 | See Table 2-39  |
| <b>MSR_UNC_CBO_3_PERFEVTSEL2</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_3_PERFEVTSEL3</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_3_UNIT_STATUS</b>                     |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_4_PERFCTR0</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_4_PERFCTR1</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_4_PERFCTR2</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |
| <b>MSR_UNC_CBO_4_PERFCTR3</b>                        |                 |
| 06_2AH .....   | See Table 2-21  |



| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                            | <b>Location</b> |
|---|-----------------|
| MSR_UNC_CBO_4_PERFEVTSELO<br>06_2AH .....                                       | See Table 2-21  |
| MSR_UNC_CBO_4_PERFEVTSEL1<br>06_2AH .....                                       | See Table 2-21  |
| MSR_UNC_CBO_4_PERFEVTSEL2<br>06_2AH .....                                       | See Table 2-21  |
| MSR_UNC_CBO_4_PERFEVTSEL3<br>06_2AH .....                                       | See Table 2-21  |
| MSR_UNC_CBO_4_UNIT_STATUS<br>06_2AH .....                                       | See Table 2-21  |
| MSR_UNC_CBO_CONFIG<br>06_2AH .....  | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....  | See Table 2-29  |
| 06_4EH, 06_5EH .....  | See Table 2-39  |
| MSR_UNC_PERF_FIXED_CTR<br>06_2AH .....  | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....  | See Table 2-29  |
| 06_4EH, 06_5EH .....  | See Table 2-39  |
| MSR_UNC_PERF_FIXED_CTRL<br>06_2AH .....   | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....  | See Table 2-29  |
| 06_4EH, 06_5EH .....  | See Table 2-39  |
| MSR_UNC_PERF_GLOBAL_CTRL<br>06_2AH .....  | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....  | See Table 2-29  |
| 06_4EH, 06_5EH .....  | See Table 2-39  |
| MSR_UNC_PERF_GLOBAL_STATUS<br>06_2AH .....                                      | See Table 2-21  |
| 06_3CH, 06_45H, 06_46H .....  | See Table 2-29  |
| 06_4EH, 06_5EH .....  | See Table 2-39  |
| MSR_UNCORE_ADDR_OPCODE_MATCH<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....    | See Table 2-15  |
| MSR_UNCORE_FIXED_CTR_CTRL<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....       | See Table 2-15  |
| MSR_UNCORE_FIXED_CTRL0<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....          | See Table 2-15  |
| MSR_UNCORE_PERF_GLOBAL_CTRL<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....     | See Table 2-15  |
| MSR_UNCORE_PERF_GLOBAL_OVF_CTRL<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH ..... | See Table 2-15  |
| MSR_UNCORE_PERF_GLOBAL_STATUS<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....   | See Table 2-15  |

## MODEL-SPECIFIC REGISTERS (MSRS)

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b>                          | <b>Location</b> |
|---|-----------------|
| MSR_UNCORE_PERFEVTSELO<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL1<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL2<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL3<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL4<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL5<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL6<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PERFEVTSEL7<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....        | See Table 2-15  |
| MSR_UNCORE_PMC0<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC1<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC2<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC3<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC4<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC5<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| 06_2EH .....  | See Table 2-16  |
| MSR_UNCORE_PMC6<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PMC7<br>06_1AH, 06_1EH, 06_1FH, 06_25H, 06_2CH .....               | See Table 2-15  |
| MSR_UNCORE_PRMRR_BASE<br>06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| MSR_UNCORE_PRMRR_MASK<br>06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| MSR_UNCORE_PRMRR_PHYS_BASE<br>06_8EH, 06_9EH .....                            | See Table 2-40  |
| MSR_UNCORE_PRMRR_PHYS_MASK<br>06_8EH, 06_9EH .....                            | See Table 2-40  |
| MSR_W_PMON_BOX_CTRL   |                 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_BOX_OVF_CTRL                              |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_BOX_STATUS                                |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_CTR0                                      |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_CTR1                                      |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_CTR2                                      |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_CTR3                                      |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_EVNT_SELO                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_EVNT_SEL1                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_EVNT_SEL2                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_EVNT_SEL3                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_FIXED_CTR                                 |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_W_PMON_FIXED_CTR_CTL                             |                 |
| 06_2EH .....   | See Table 2-16  |
| MSR_WEIGHTED_CORE_CO                                 |                 |
| 06_4EH, 06_5EH, 06_55H, 06_8EH, 06_9EH, 06_66H ..... | See Table 2-38  |
| MTRRfix16K_80000                                     |                 |
| 06_0EH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| MTRRfix16K_A0000                                     |                 |
| 06_0EH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| MTRRfix4K_C0000                                      |                 |
| 06_0EH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| MTRRfix4K_C8000                                      |                 |
| 06_0EH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| MTRRfix4K_D0000                                      |                 |
| 06_0EH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |

| MSR Name and CPUID DisplayFamily_DisplayModel | Location       |
|---|----------------|
| MTRRfix4K_D8000                               |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRfix4K_E0000                               |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRfix4K_E8000                               |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRfix4K_F0000                               |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRfix4K_F8000                               |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRfix64K_00000                              |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase0                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase1                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase2                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase3                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase4                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase5                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase6                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |
| MTRRphysBase7                                 |                |
| 06_OEH .....                                  | See Table 2-48 |
| P6 Family .....                               | See Table 2-50 |

| <b>MSR Name and CPUID DisplayFamily_DisplayModel</b> | <b>Location</b> |
|--|-----------------|
| <b>MTRRphysMask0</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask1</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask2</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask3</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask4</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask5</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask6</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |
| <b>MTRRphysMask7</b>                                 |                 |
| 06_OEH .....   | See Table 2-48  |
| P6 Family .....                                      | See Table 2-50  |

