

Intel® Itanium® Architecture Software Developer's Manual Specification Update

October 2009



THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://developer.intel.com/design/litcentr>.

Intel and Itanium are trademarks of Intel Corporation in the U. S. and other countries.

Copyright © 2002-2009, Intel Corporation. All rights reserved.

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



Contents

1	Preface	5
2	Summary Table of Changes	6
3	Specification Changes.....	7
4	Specification Clarifications	59
5	Documentation Changes	66



Revision History

Document Number	Version Number	Description	Date
248699	-013	Specification change 20.	October 2009
248699	-012	Specification changes 15 - 19.	June 2008
248699	-011	Specification changes 1-14, Specification Clarifications 1-10, and Document Change 1.	October 2007
248699	-001 - -010	Changes from previous Software Developer's Manual Specification Updates were incorporated into version 2.2 of the <i>Intel® Itanium® Architecture Software Developer's Manual</i> January 2006.	January 2006

§



1 Preface

This document is an update to the specifications contained in the Affected Documents/Related 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.

Information types defined in Nomenclature are consolidated into the specification update and are no longer published in other documents.

This document may also contain information that was not previously published.

1.1 Affected Documents/Related Documents


Title	Document #
<i>Intel® Itanium® Architecture Software Developer's Manual, Volume 1: Application Architecture, Revision 2.2</i>	245317-005
<i>Intel® Itanium® Architecture Software Developer's Manual, Volume 2: System Architecture, Revision 2.2</i>	245318-005
<i>Intel® Itanium® Architecture Software Developer's Manual, Volume 3: Instruction Set Reference, Revision 2.2</i>	245319-005

1.2 Nomenclature

Specification Changes are modifications to the current published specifications for Intel® Itanium® processors. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further explain a specification's interpretation. These clarifications will be incorporated in the next release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the *Intel® Itanium® Architecture Software Developer's Manual*.



2 Summary Table of Changes

The following tables indicate the specification changes and specification clarifications that apply to the *Intel® Itanium® Architecture Software Developer's Manual*.

2.1 Specification Changes

No.	Page	SPECIFICATION CHANGES
1	8	Illegal VAC/VDC combinations and IIPA requirements
2	9	Resource Utilization Counter
3	20	PAL_VP_INIT and VPD.vpr changes
4	19	New PAL_VPS_RESUME_HANDLER to indicate RSE Current Frame Load Enable setting at the target instruction
5	20	PAL_VP_INIT_ENV Implementation-specific Configuration Option
6	20	Increase in minimum number of virtual address bits
7	21	PAL_MC_ERROR_INFO health indicator
8	24	New implementation-specific bit fields for PAL_MC_ERROR_INJECT
9	24	MOV-to-PSR.L Reserved Field Checking
10	24	Virtual Machine Disable
11	27	Removal of pal_proc_vector argument from PAL_VP_SAVE and PAL_VP_RESTORE
12	27	Variable Frequency Mode Additions to ACPI P-states
13	30	PAL_MC_DYNAMIC_STATE Changes
14	32	Min-State Save Area Size Change
15	33	Data Speculation Disable
16	34	Interrupt Instruction Bundle Registers
17	44	Data-Poisoning Promotion Changes
18	45	ACPI P-State Clarifications
19	54	Synchronization Requirements for Virtualization Opcode Optimization
20	58	New Priority Hint and Recommendations for Multi-Threading Hints

2.2 Specification Clarifications

No.	Page	SPECIFICATION CLARIFICATIONS
1	62	Clarification of ptc.g release semantics
2	62	Clarification of PAL_MC_ERROR_INFO reporting of uncacheable transactions
3	63	Clarification of behavior when ptc.g overlaps a translation register
4	63	INT3 Clarifications
5	63	Test feature instruction clarifications
6	63	Clarification of performance counter behavior under halt states
7	65	PMI Clarifications
8	66	PAL_MC_ERROR_INJECT Clarifications
9	67	Min-state Save Area Clarifications
10	67	Semaphore Code Corrections

2.3 Documentation Changes

No.	Page	DOCUMENTATION CHANGES
1	69	Revision 2.2 Documentation Changes

3 Specification Changes

1. Illegal VAC/VDC combinations and IIPA requirements

1. In Volume 2, Part I, Chapter 11, add a new Section 11.7.4.3 “Virtualization Optimization Combinations”. In the new Section 11.7.4.3, add the following description and table to indicate the allowed *vac* and *vdc* combinations:

11.7.4.3 Virtualization Optimization Combinations

Table 11-35 describes the supported combinations of virtualization accelerations and disables.

Table 11-35. Supported Virtualization Optimization Combinations

	d_vmsw	d_extint	d_ibr_dbr	d_pmc	d_to_pmd	d_itm	d_psr_i
a_int	o ^a	x ^b	o	o	o	o	x
a_from_int_cr	o	o	o	o	o	o	o
a_to_int_cr	o	o	o	o	o	o	o
a_from_psr	o	o	o	o	o	o	x
a_from_cpuid	o	o	o	o	o	o	o
a_cover	o	o	o	o	o	o	o
a_bsw	o	o	o	o	o	o	x

Notes:

- a. “o” indicates the corresponding virtualization acceleration and disable can be enabled together.
- b. “x” indicates the corresponding virtualization acceleration and disable cannot be enabled together.

2. In Volume 2, Part I, Chapter 11, Section 11.7.4.2.7, “Disable PSR Interrupt-bit Virtualization”, replace the note:

Note: This field overrides the a_int Virtualization Acceleration Control (*vac*) described in Section 11.7.4.1.1, “Virtual External Interrupt Optimization” on page 2:323. If this control is enabled (set to 1), the a_int Virtualization Acceleration Control (*vac*) is ignored.

with:

Note: This field cannot be enabled together with a_int, a_from_psr or a_bsw virtualization accelerations. If this control is enabled together with any one of the described accelerations, an error will be returned during PAL_VP_CREATE and PAL_VP_REGISTER. See Section 11.7.4.3 for details.

3. In Volume 2, Part I, Chapter 11, Section 11.7.4.2.2 “Disable External Interrupt Control Register Virtualization”, replace the note:

Note: This field overrides the a_int Virtualization Acceleration Control (*vac*) described in Section 11.7.4.1.1, “Virtual External Interrupt Optimization” on page 2:323. If this control is enabled (set to 1), the a_int Virtualization Acceleration Control (*vac*) is ignored.

with:

Note: This field cannot be enabled together with the a_int virtualization acceleration control (*vac*) described in Section 11.7.4.1.1. If this control is enabled together with the a_int control, an error will be returned during PAL_VP_CREATE and PAL_VP_REGISTER. See Section 11.7.4.3 for details.



4. In Volume 2, Part I, Chapter 11, Section 11.10, “PAL Procedures”, PAL_VP_CREATE page. In the Description section, after the second paragraph, last sentence:
The *vac*, *vdc* and *virt_env_vaddr* parameters in the VPD must already be initialized before calling this procedure.

Add the following sentence:
Invalid argument is returned on unsupported *vac/vdc* combinations. See Section 11.7.4.3 for details.
5. In Volume 2, Part I, Chapter 11, Section 11.10, “PAL Procedures”, PAL_VP_REGISTER page. In the Description section, add the following paragraph before the last paragraph:
PAL_VP_REGISTER returns invalid argument on unsupported virtualization optimization combinations in *vpd*. See Section 11.7.4.3, “Virtualization Optimization Combinations” for details.
6. In Volume 2, Part I, Chapter 11, Section 11.7.4.1.1 “Virtual External Interrupt Optimization”, add the following note to the end of the section:
Note: This field cannot be enabled together with *d_extint* or *d_psr_i* virtualization disables. If this control is enabled together with any one of described disables, an error will be returned during PAL_VP_CREATE and PAL_VP_REGISTER. See Section 11.7.4.3, “Virtualization Optimization Combinations” for details.
7. In Volume 2, Part I, Chapter 11, Section 11.7.4.1.4 “MOV-from-PSR Optimization”, add the following note to the end of the section:
Note: This field cannot be enabled together with the *d_psr_i* virtualization disable control (*vdc*) described in Section 11.7.4.2.7, “Disable PSR Interrupt-bit Virtualization”. If this control is enabled together with the *d_psr_i* control, an error will be returned during PAL_VP_CREATE and PAL_VP_REGISTER. See Section 11.7.4.3, “Virtualization Optimization Combinations” for details.
8. In Volume 2, Part I, Chapter 11, Section 11.7.4.1.7 “Bank Switch Optimization”, add the following note to the end of the section:
Note: This field cannot be enabled together with the *d_psr_i* virtualization disable control (*vdc*) described in Section 11.7.4.2.7. If this control is enabled together with the *d_psr_i* control, an error will be returned during PAL_VP_CREATE and PAL_VP_REGISTER. See Section 11.7.4.3, “Virtualization Optimization Combinations” for details.

2. Resource Utilization Counter

The existing Interval Time Counter application register is clocked at a constant rate, independent of logical processor and virtual processor context switches on a processor core.

The new Resource Utilization Counter application register is clocked like the ITC, but is provided per logical or virtual processor and provides an estimate of the portion of resources used by a logical or virtual processor with respect to all resources provided by the underlying physical processor.

1. Add AR.ruc to Volume 1, Part I, Section 3.1.1, “Reserved and Ignored Registers and Fields”. In Figure 3-1, add RUC, AR 45 directly below and directly next to ITC, AR 44.



2. In Volume I, Part I, Section 3.1.8, “Application Registers”, Table 3-3, add the RUC entry as shown:

Register	Name	Description	Execution Unit Type
AR 44	ITC	Interval Time Counter	
AR 45	RUC	Resource Utilization Counter	
AR 46 - AR 47		Reserved	

- a. Add a new Resource Utilization Counter section just after Section 3.1.8.10, “Interval Time Counter”:

3.1.8.11 Resource Utilization Counter (RUC - AR 45)

The Resource Utilization Counter (RUC) is a 64-bit register which provides an estimate of the portion of resources used by a logical or virtual processor with respect to all resources provided by the underlying physical processor.

In a given time interval, the difference in the RUC values for all of the logical or virtual processors on a given physical processor add up to approximately the difference seen in the ITC on that physical processor for that same interval. (See Vol 2, Section 11.7 for details on virtual processors.)

A sequence of reads of the RUC is guaranteed to return ever-increasing values (except for the case of the counter wrapping back to 0) corresponding to the program order of the reads.

System software can secure the resource utilization counter from non-privileged access. When secured, a read of the RUC at any privilege level other than the most privileged causes a Privileged Register fault.

3. Add CPUID[4] bit for non-privileged discovery of the Resource Utilization Counter
 - a. In Volume 1, Part I, section 3.1.11, “Processor Identification Registers”, Figure 3-12, add a new *ru* bit as shown:

Figure 3-12. CPUID Register 4 – General Features/Capability Bits



- b. In Volume 1, Part I, Section 3.1.11, “Processor Identification Registers”, add the *ru* entry to Table 3-8 as shown:

Table 3-8. CPUID Register 4 Fields

Field	Bits	Description
ru	3	Processor implements the Resource Utilization Counter (AR 45).
rv	63:4	Reserved.

4. In Volume 1, Part I, Section 6.2, “IA-32 Application Register State Model”, add AR.ruc to Fig 6-3 and Table 6-1.
 - a. In Figure 6-3, add RUC, AR 45 directly below and directly next to ITC, AR 44.
 - b. In Table 6-1, add a new row, just under the row for ITC, with this information:

RUC		Unmodified	64	RUC continues to count while in IA-32 execution mode
-----	--	------------	----	--



5. Describe serialization for RUC in Volume 2, Part I, Section 3.2.2 “Data Serialization”.
 - a. In the second paragraph, change:
... RR, PKR, DTR, ...
to:
... RR, PKR, RUC, DTR, ...
6. Add AR.ruc to Volume 2, Part I, Section 3.3.1 “System State Overview”.
 - a. In the bullet list, just after the bullet for “Interval Timer Facilities”, add a new bullet:
 - Resource Utilization Facility - A 64-bit resource utilization counter is provided for privileged and non-privileged use. This counts the number of Interval Timer cycles consumed by this logical processor. See Section 3.1.8.11, “Resource Utilization Counter” on page 1:29.
 - b. Add AR.ruc to Figure 3-1:
Add RUC, AR 45 directly below and directly next to ITC, AR 44.
7. Update the description of PSR.si to cover RUC.

In Volume 2, Part I, Section 3.3.2 “Processor Status Register (PSR)”, Table 3-2, change this part of the description of the PSR.si bit:

When 1, the Interval Time Counter (ITC) register is readable only by privileged code; non-privileged reads result in a Privileged Register fault. When 0, ITC is readable at any privilege level.

to:


When 1, the Interval Time Counter (ITC) register and the Resource Utilization Counter (RUC) are readable only by privileged code; non-privileged reads result in a Privileged Register fault. When 0, ITC and RUC are readable at any privilege level.
8. Add AR.ruc to Volume 2, Part I, Section 3.3.4 “Global Control Registers”.
 - a. Add a new section, just after 3.3.4.2 “Interval Time Counter and Match Register”:

3.3.4.3 Resource Utilization Counter (AR 45)

The Resource Utilization Counter (RUC) is a 64-bit counter that counts up at a fixed relationship to the input clock to the processor, when the processor is active. Processors may be inactive due to hardware multi-threading. Virtual processors may be inactive when not scheduled to run by the VMM. (See Vol 2, section 11.7 for details on virtual processors.)

The RUC is clocked such that, in a given time interval, the difference in the RUC values for all of the logical or virtual processors on a given physical processor add up to approximately the difference seen in the ITC on that physical processor for that same interval.

A sequence of reads of the RUC is guaranteed to return ever-increasing values (except for the case of the counter wrapping back to 0) corresponding to the program order of the reads. Applications can directly sample the RUC for active-running-time calculations.



A 64-bit overflow condition can occur without notification. The RUC can be read at any privilege level if PSR.si is zero. The timer can be secured from non-privileged access by setting PSR.si to one. When secured, a read of the RUC by non-privileged code results in a Privileged Register fault. Writes to the RUC can only be performed at privilege level 0; otherwise, a Privileged Register fault is raised.

Modification of the RUC is not necessarily serialized with respect to instruction execution. Software can issue a data serialization operation to ensure the RUC updates are observed by a given point in program execution. Software must accept a level of sampling error when reading the resource utilization counter due to various machine stall conditions, interruptions, bus contention effects, and so forth. Please see the processor-specific documentation for further information on the level of sampling error of the Itanium processor.

RUC should only be written by Virtual Machine Monitors; other Operating Systems should not write to RUC, but should only read it.

9. Update Volume 2, Part I, Section 3.4, "Processor Virtualization", Table 3-10. Change the last two rows to the following:

Class	Virtualized Instructions
Reading AR[ITC] or AR[RUC] with PSR.si == 1 (virtualized at all privilege levels)	mov from ar.itc mov from ar.ruc
Instructions which write privileged registers	mov to itc mov to ar.ruc

10. Update Volume 2, Part I, Section 11.3.2, "PALE_CHECK Exit State".

- a. In the bullet for ARs, change the following:
The contents of all application registers are unchanged from the time of the MCA, except the RSE control register (RSC), the RSE backing store pointer (BSP), and the ITC counter.
to:
The contents of all application registers are unchanged from the time of the MCA, except the RSE control register (RSC), the RSE backing store pointer (BSP), and the ITC and RUC counters.
- b. In that same AR bullet, change the following:
The ITC register will not be directly modified by PAL, but will continue to count during the execution of the MCA handler.
to:
The ITC register will not be directly modified by PAL, but will continue to count during the execution of the MCA handler. The RUC register will not be directly modified by PAL, but will Continue to count during the execution of the MCA handler while the processor is active.

11. Update Volume 2, Part I, Section 11.4.2, "PALE_INIT Exit State".

- a. In the bullet for ARs, change the following:
The contents of all application registers are unchanged from the time of the INIT, except the RSE control register (RSC), the RSE backing store pointer (BSP), and the ITC counter.
to:
The contents of all application registers are unchanged from the time of the INIT, except the RSE control register (RSC), the RSE backing store pointer (BSP), and the ITC and RUC counters.



- b. In that same AR bullet, change the following:
 The ITC register will not be directly modified by PAL, but will continue to count during the execution of the INIT handler.
 to:
 The ITC register will not be directly modified by PAL, but will continue to count during the execution of the INIT handler. The RUC register will not be directly modified by PAL, but will continue to count during the execution of the INIT handler while the processor is active.

12. Update Volume 2, Part I, Section 11.5.2, "PALE_PMI Exit State".

- a. In the bullet for ARs, change the following:
 The contents of all application registers are unchanged from the time of the interruption, except the RSE control register (RSC) and the ITC counter.
 to:
 The contents of all application registers are unchanged from the time of the interruption, except the RSE control register (RSC) and the ITC and RUC counters.
- b. In that same AR bullet, change the following:
 The ITC register will not be directly modified by PAL, but will continue to count during the execution of the PMI handler.
 to:
 The ITC register will not be directly modified by PAL, but will continue to count during the execution of the PMI handler. The RUC register will not be directly modified by PAL, but will continue to count during the execution of the PMI handler while the processor is active.

13. Update Volume 2, Part I, Section 11.7.3.1, "PAL Virtualization Intercept Handoff State".

- a. In the bullet for ARs, change the following:
 The contents of all application registers are preserved from the time of the interruption, except the ITC counter. The ITC register will not be directly modified by PAL, but will continue to count during the execution of the virtualization intercept handler.
 to:
 The contents of all application registers are preserved from the time of the interruption, except the ITC and RUC counters. The ITC register will not be directly modified by PAL, but will continue to count during the execution of the virtualization intercept handler. The RUC register will not be directly modified by PAL, but will continue to count during the execution of the virtualization intercept handler while the processor is active.

14. Update Volume 2, Part I, Section 11.10.2.2.7, "Application Registers". Add a new row to Table 11-48, just below the ITC row, with this information:

Register	Description	Class
RUC	Resource Utilization Counter	unchanged ^c

c. No PAL procedure writes to the RUC. The value at exit is the value at entry plus the number of cycles provided to the processor during the procedure call.



15. Update Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_ENTER_IA_32_ENV page.

a. Add a new row to the bottom of Table 11-67:


Intel® Itanium® Register	IA-32 State	Description
RUC	--	RUC continues to count while in IA-32 execution mode

b. Add a new row to Table 11-71 just below the row for ITC:

Intel® Itanium® Register	IA-32 State	Description
RUC	--	Final value of RUC

16. Add a new PAL_VP_INFO procedure for privileged discovery of the Resource Utilization Counter.

a. Add a new PAL procedure to Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", just before PAL_VP_INIT_ENV:



PAL_VP_INFO – PAL Virtual Processor Information (50)

Purpose: Returns information about virtual processor features.

Calling Conv: Static


Mode: Physical

Buffer: Not dependent

Arguments:	Argument	Description
	index	Index of PAL_VP_INFO within the list of PAL procedures.
	feature_set	Feature set information is being requested for.
	vp_buffer	An address to an 8-byte aligned memory buffer (if used).
	Reserved	0.
Returns:	Return Value	Description
	status	Return status of the PAL_VP_INFO procedure.
	vp_info	Information about the virtual processor..
	vmm_id	Unique identifier for the VMM.
	Reserved	0
Status:	Status Value	Description
	0	Call completed without error
	-1	Unimplemented procedure
	-2	Invalid argument
	-3	Call completed with error
	-8	Specified <i>feature_set</i> is not implemented

Description: The PAL_VP_INFO procedure call is used to describe virtual processor features.

The *feature_set* input argument for PAL_VP_INFO describes which virtual-processor *feature_set* information is being requested, and is composed of two fields as shown:




A *vmm_id* of 0 indicates architected feature sets, while others are implementation-specific feature sets. Implementation-specific feature sets are described in VMM-specific documentation.

This procedure will return a -8 if an unsupported *feature_set* argument is passed as an input. The return status is used by the caller to know which feature sets are currently supported on a particular VMM. This procedure always returns unimplemented (-1) when called on physical processors.

For each valid *feature_set*, this procedure returns information about the virtual processor in *vp_info*. Additional information may be returned in the memory buffer pointed to by *vp_buffer*, as needed. Details, for a given implementation-specific *feature_set*, of whether information is returned in the buffer, the size of the buffer, and the representation of this information in the buffer and in *vp_info* are described in VMM-specific documentation.

Architected *feature_set* 0 (*vmm_id* 0, *index* 0) is defined and required to be implemented (if this procedure is implemented), but there are no architected features defined in it yet, and so all bits in *vp_info* are reserved for architected *feature_set* 0.



Other architected feature sets (*vmx_id* 0, *index*>0) are undefined, and return -8 (Specified *feature_set* is not implemented). SW can call PAL_VP_INFO with a *feature_set* argument of 0 to get the *vmx_id*, although *vmx_id* is also returned for any other implemented feature sets as well. For *feature_set* 0, the *vp_buffer* arg is ignored.

- b. Add PAL_VP_INFO to Volume 2, Part I, Section 11.10.1, "PAL Procedure Summary", Table 11-42.

Procedure	Idx	Class	Conv.	Mode	Buffer	Description
PAL_VP_INFO	50	Opt.	Static	Phys.	No.	Returns information about virtual processor features.

17. In Volume 3, Part I, Section 2.2, "Instruction Descriptions", update the "mov ar" instruction page pseudo-code.

- a. In the Operation section, change the following "from form" code:

```
if (ar3 == ITC && PSR.si && PSR.cpl != 0)
    privileged_register_fault();

if (ar3 == ITC && PSR.si && PSR.vm == 1)
    virtualization_fault();"
to:
if ((ar3 == ITC || ar3 == RUC) && PSR.si && PSR.cpl != 0)
    privileged_register_fault();

if ((ar3 == ITC || ar3 == RUC) && PSR.si && PSR.vm == 1)
    virtualization_fault();"
```

- b. In the to_form code, change the following "to form" code:

```
if ((is_kernel_reg(ar3) || ar3 == ITC) && (PSR.cpl != 0))
    privileged_register_fault();

if (ar3 == ITC && PSR.vm == 1)
    virtualization_fault();
to:
if ((is_kernel_reg(ar3) || ar3 == ITC || ar3 == RUC) && (PSR.cpl != 0))
    privileged_register_fault();

if ((ar3 == ITC || ar3 == RUC) && PSR.vm == 1)
    virtualization_fault();
```

- c. At the beginning of the Operation section, change the following code:

```
if (is_reserved_reg(tmp_type, ar3))
    illegal_operation_fault();
to:
if (!instruction_implemented(MOV_AR_RUC))
    illegal_operation_fault();

if (is_reserved_reg(tmp_type, ar3))
    illegal_operation_fault();
```

18. In Volume 3, Part I, Section 5.3.2, "RAW Dependency Table", add AR.ruc to the resource dependency tables.

- a. In Table 5-2, add a row just under the row for AR[RSC] with this information:

Resource Name	Writers	Readers	Semantics of Dependency
AR[RUC]	mov-to-AR-RUC	br.ia, mov-from-AR-RUC	impliedF



- b. In Table 5-2, in the row for PSR.cpl, add mov-from-AR-RUC and mov-to-AR-RUC to the readers in both of the sub-rows as shown:

Resource Name	Writers	Readers	Semantics of Dependency
PSR.cpl	epc, br.ret	priv-ops , br.call, brl.call, epc, PAL_VP_INIT_ENV Implementation-specific Configuration Option , mov-to-AR-ITC, mov-to-AR-RSC, mov-to-AR-K, mov-from-IND-PMD, probe-all, mem-readers, mem-writers, lfetch-all, mov-from-AR-RUC, mov-to-AR-RUC	implied
	rfi	priv-ops , br.call, brl.call, epc, mov-from-AR-ITC, mov-to-AR-ITC, mov-to-AR-RSC, mov-to-AR-K, mov-from-IND-PMD, probe-all, mem-readers, mem-writers, lfetch-all, mov-from-AR-RUC, mov-to-AR-RUC	impliedF

- c. In Table 5-2, update the PSR.si row, as shown:

Resource Name	Writers	Readers	Semantics of Dependency
PSR.si	sys-mask-writers-partial ⁷ , mov-to-PSR-I	mov-from-PSR	impliedF
		mov-from-AR-ITC, mov-from-AR-RUC,	data
	rfi	mov-from-AR-ITC, mov-from-AR-RUC, mov-from-PSR	impliedF

- d. In Table 5-2, update the row for PSR.vm as shown:

Resource Name	Writers	Readers	Semantics of Dependency
PSR.vm	vmsw	mem-readers , mem-writers , mov-from-AR-ITC , mov-from-IND-CPUID , mov-to-AR-ITC , priv-ops \vmsw, cover, thash, ttag, mov-from-AR-RUC , mov-to-AR-RUC	implied
	rfi	mem-readers , mem-writers , mov-from-AR-ITC , mov-from-IND-CPUID , mov-to-AR-ITC , priv-ops \vmsw, cover, thash, ttag, mov-from-AR-RUC, mov-to-AR-RUC	impliedF

- e. In Volume 2, Part I, Section 5.3.3, “WAW Dependency Table”, Table 5-3, add a row just under the row for AR[RSC] with this information:

Resource Name	Writers	Semantics of Dependency
AR[RUC]	mov-to-AR-RUC	impliedF



- f. In Volume 3, Part I, Section 5.4 “Support Tables”, Table 5-5, add a row just under the row for mov-from-AR-RSC with this information:

Class	Events/Instructions
mov-from-AR-RUC	mov-from-AR-M[Field(ar3) == RUC]

- g. In Volume 3, Part I, Section 5.4 “Support Tables”, Table 5-5, add a row just under the row for mov-to-AR-RSC with this information:

Class	Events/Instructions
mov-to-AR-RUC	mov-to-AR-M[Field(ar3) == RUC]

19. In Volume 2, Part II, add a new section just after Section 10.5.5, “Interval Timer Usage Example”:

10.5.6 Resource Utilization Counter Usage Example

The Itanium architecture provides a 64-bit counter to provide information on how many execution cycles a given logical processor is getting. It is similar to the Interval Timer (ITC, AR 44), except that it is clocked only when the logical processor is active. Optimizations such as hardware multi-threading and processor virtualization may cause a logical processor to sometimes be inactive. The Resource Utilization Counter allows for better cycle accounting for logical processors, given these types of optimizations. RUC should only be written by Virtual Machine Monitors; other Operating Systems should not write to RUC, but should only read it.

3. PAL_VP_INIT and VPD.vpr Changes

1. PAL_VP_INIT_ENV currently freezes performance registers by clearing PMC[0].fr when fr_pmc is enabled. The following change removes the race condition that can occur when a counter overflow happens just before the write to PMC[0].fr, causing PAL to overwrite the overflow bit, and losing overflow information.

- a. In Volume 2, Part I, Section 11.10.3, “PAL Procedure Specifications”, PAL_VP_INIT_ENV procedure, Table 11-110, change the *fr_pmc* description from:

fr_pmc	1	If 1, performance counters are frozen on all IVA-based interruptions when virtual processors are running. If 0, the performance counters will not be frozen on IVA-based interruptions when virtual processors are running.
--------	---	---

to:


fr_pmc	1	If 1, for virtualization intercepts the performance counters are disabled by setting PSR.up and pp to 0, see Section 11.7.3.1 for details on PSR settings at virtualization intercepts; for all other IVA-based interruptions PSR.pp and up are set according to Interruption State column described in Processor Status Field table described in Vol 2 Table 3-2. If 0, PSR.pp and up are set according to Interruption State column described in Processor Status Field table described in Vol 2 Table 3-2.
--------	---	---

- b. Volume 2, Part I, Section 11.7.3.1, “PAL Virtualization Handoff State”, PSR description. Change the following bullet:

- PSR: PSR fields are set according to the “Interruption State” column in Table 3-2, “Processor Status Register Fields” on page 2:21.

to:

- PSR: PSR fields are set according to the “Interruption State” column in Table 3-2, “Processor Status Register Fields” on page 2:21. PSR.up and pp



are set to 0 when *fr_pmc* field in *config_options* parameter during PAL_VP_INIT_ENV is 1.

2. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_VP_INIT_ENV page, Table 11-110, change the *be config_options* description from:

be	2	Big-endian – Indicates the endian setting of the VMM. If 1, the values in the VPD are stored in big-endian format and the PAL services calls are made with PSR.be bit equals to 1. If 0, the values in the VPD are stored in little-endian format and the PAL services calls are made with PSR.be bit equals to 0.
----	---	--

to:

be	2	Big-endian – Indicates the endian setting of the VMM. If 1, the values in the VPD are stored in big-endian format and the PAL services calls are made with PSR.be bit equals to 1. If 0, the values in the VPD are stored in little-endian format and the PAL services calls are made with PSR.be bit equals to 0. The VMM must match DCR.be with the value set in this field when the IVA control register on the logical processor is set to point to the per-virtual-processor host IVT. See Section 11.17.2 "Interrupt Handling in a Virtual Environment" and Table 11-17 "IVA Settings after PAL Virtualization-related Procedures and Services" for details on per-virtual-processor host IVT.
----	---	--

3. Volume 2, Part I, Section 11.7.1, "Virtual Processor Descriptor", Table 11-14, Update the *vpr* row from:

vpr	1	1432	Virtual Predicate Registers – Represents the Predicate Registers of the virtual processor. The bit positions in vpr correspond to predicate registers in the same manner as with the mov predicates instruction.	Architectural State [always]
-----	---	------	--	------------------------------

to:

vpr	1	1432	Virtual Predicate Registers – Represents the Predicate Registers of the virtual processor. The bit positions in vpr correspond to predicate registers in the same manner as with the mov predicates instruction. The contents in this field are undefined except at virtualization intercept handoff. The VMM can not rely on the contents in this field to be preserved when the virtual processor is running.	Architectural State [always]
-----	---	------	---	------------------------------

4. **New PAL_VPS_RESUME_HANDLER to Indicate RSE Current Frame Load Enable Setting at the Target Instruction**

This change adds a parameter to PAL_VPS_RESUME_HANDLER to allow a virtual machine monitor to specify the register stack engine Current Frame Load Enable setting at the target guest handler.


1. Volume 2, Part I, Section 11.11.12, "PAL Virtualization Service Specifications", PAL_VPS_RESUME_HANDLER page.
 - a. Change the description of GR26 from:

Virtualization Acceleration Control (*vac*) field from the VPD specified in GR25

to:

Virtualization Acceleration Control (*vac*) field from the VPD specified in GR25 and CFLE setting at the target instruction.
 - b. In the second paragraph, change the following paragraph from:

The VMM specifies the BR0 of the virtual processor in GR24, the 64-bit virtual pointer to the VPD in GR25 and the *vac* field of the VPD in GR26.



Behavior is undefined if the *vac* in GR26 does not match the *vac* field in the VPD argument specified in GR25.

to:

GR24 specifies the BR0 of the virtual processor; GR25 specifies the 64-bit virtual pointer to the VPD; GR26 specifies the *vac* field of the VPD in GR26; bit 63 of GR26 specifies the value of CFLE setting at the target instruction. Behavior is undefined if the *vac* in GR26 does not match the *vac* field in the VPD argument specified in GR25.

2. In Volume 3, Part I, Section 2.2, “Instruction Descriptions”, “tf” page:
 - a. In the Description section, remove the following text:
Implementation of PSR.vm is optional. If it is implemented but the instruction is disabled, this instruction takes Virtualization fault when executed with PSR.vm equals to 1.
 - b. In the Operation section, remove both instances of the following lines:

```
if (PSR.vm == 1 && vm_tf_disabled())  
virtualization_fault();
```
 - c. In the Interruptions section, remove “Virtualization fault” from the list of interruptions.

5. PAL_VP_INIT_ENV Implementation-specific Configuration Option

This change defines an implementation-specific configuration bit for PAL_VP_INIT_ENV. A separate update to the *Dual-Core Update to the Itanium 2 Processor Reference Manual* will define this implementation-specific bit to optimize performance for virtual machine monitors using data translation cache for pages containing virtualized instructions.

1. Volume 2, Part 1, Section 11.10.3, “PAL Procedures Specifications”, PAL_VP_INIT_ENV page. Add the *impl* bit to Table 11-110:

Field	Bit	Description
impl	63	Implementation-specific configuration option. This field is ignored if not implemented. Please refer to processor-specific documentation for details.

2. A future revision of the *Dual-Core Update to the Itanium 2 Processor Reference Manual* will include the following change:

Add a section to describe PAL_VP_INIT_ENV:

Table 11-62. PAL_VP_INIT_ENV Implementation-specific Behavior

Field	Bit	Description
hint_dtc	63	If 1, this hint indicates the VMM is using data translation cache for pages containing virtualized instructions. Instruction TLB misses will happen during virtualized instruction execution if the corresponding data translation does not exist in the TLB hierarchy

6. Increase in Minimum Number of Virtual Address Bits

This change increases the minimum number of implemented virtual address bits from 51 to 54. Note that Itanium 2 processors and Dual Core Itanium 2 processors already support the 54 bit virtual address minimum.

1. Volume 2, Part 1, Section 4.3.2, “Unimplemented Virtual Address Bits” change the first paragraph:



... all processor models implement at least 51 virtual address bits; i.e., the smallest IMPL_VA_MSB is 50.

to:

... all processor models implement at least 54 virtual address bits; i.e., the smallest IMPL_VA_MSB is 53.

2. Change the second paragraph of Volume 2, Part 1, Section 4.3.2, "Unimplemented Virtual Address Bits" from:

If the PSR.vm bit is implemented, at least 52 virtual address bits must be implemented.

to:

If the PSR.vm bit is implemented, at least 55 virtual address bits must be implemented.

7. PAL_MC_ERROR_INFO Health Indicator

This change defines PAL_MC_ERROR_INFO cache_check, tlb_check, and uarch_check fields to allow hardware status tracking to be reported for processor structures. A new PAL_MC_HW_TRACKING procedure allows software to determine which processor structures provide hardware status tracking.

1. In Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_MC_ERROR_INFO page, Table 11-82, add a *hlth* field row and a table note to the cache_check definition:

Field	Bits	Description
hlth	31:30	Health indicator. This field will report if the cache type and level reporting this error supports hardware status tracking and the current status of this cache. 00 - No hardware status tracking is provided for the cache type and level reporting this event. 01 - Status tracking is provided for this cache type and level and the current status is normal status. ^(a) 10 - Status tracking is provided for the cache type and level and the current status is cautionary. ^(a) When a cache reports a cautionary status the "hardware damage" bit of the PSP (See Section 11.3.2.1, "Processor State Parameter (GR18)") will be set as well. 11 - Reserved
rsvd	31:24	Reserved

(a) Hardware is tracking the operating status of the structure type and level reporting the error. The hardware reports a "normal" status when the number of entries within a structure reporting repeated corrections is at or below a pre-defined threshold. A "cautionary" status is reported when the number of affected entries exceeds a pre-defined threshold.

2. In Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_MC_ERROR_INFO page, Table 11-83, add a new *hlth* field row and a table note to the tlb_check definition:

Field	Bits	Description
hlth	31:30	Health indicator. This field will report if the tlb type and level reporting this error supports hardware status tracking and the current status of this tlb. 00 - No hardware status tracking is provided for the tlb type and level reporting this event. 01 - Status tracking is provided for this tlb type and level and the current status is normal. ^(a) 10 - Status tracking is provided for the tlb type and level and the current status is cautionary. ^(a) When a tlb reports a cautionary status the "hardware damage" bit of the PSP (See Section 11.3.2.1, "Processor State Parameter (GR18)") will be set as well. 11 - Reserved
rsvd	31:24	Reserved

(a) Hardware is tracking the operating status of the structure type and level reporting the error. The hardware reports a "normal" status when the number of entries within a structure reporting repeated corrections is at or below a pre-defined threshold. A "cautionary" status is reported when the number of affected entries exceeds a pre-defined threshold.



3. Add a new PAL procedure called PAL_MC_HW_TRACKING

- a. Add PAL_MC_HW_TRACKING to Volume 2, Part I, Section 11.10.1, "PAL Procedure Summary", Table 11-38:

Procedure	Idx	Class	Conv.	Mode	Buffer	Description
PAL_MC_HW_TRACKING	51	Opt.	Static	Both	Yes	Query which hardware structures are performing hardware status tracking.

- b. In Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", add a new PAL_MC_HW_TRACKING page after the PAL_MC_EXPECTED page:




Buffer: Not dependent

The layout of the tracked return value is showing in Fig 11-64.

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
Reserved																DTT			ITT			DCT			ICT						
63	62	61	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32
Reserved																															

Field	Bits	Description
ICT	3:0	Instruction cache tracking. This is a 4-bit vector denoting which levels of instruction cache provide hardware tracking.
DCT	7:4	Data cache tracking. This is a 4-bit vector denoting which levels of the data/unified caches provide hardware tracking.
ITT	11:8	Instruction TLB tracking. This is a 4-bit vector denoting which levels of the instruction TLB provide hardware tracking.
DTT	15:12	Data TLB tracking. This is a 4-bit vector denoting which levels of data/unified TLB provide hardware tracking
rsvd	63:16	Reserved.



The convention for the levels in the *hw_track* field is such that the least significant bit in the field represents the lowest level of the structures hierarchy. For example bit 0 of the ICT field represents the first level instruction cache.

8. New Implementation-specific Bit Fields for PAL_MC_ERROR_INJECT

This change defines implementation-specific bits for PAL_MC_ERROR_INJECT.

1. In Volume 2, Part I, Section 11.10.3, PAL_MC_ERROR_INJECT procedure, Figure 11-51 change bits 63:48 from Reserved to implementation specific in Figure 11-51 as shown:

Figure 11-51. *err_type_info*

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
Reserved																struct_hier			err_struct			err_sev			err_inj			mode			
63	62	61	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32
Impl_Spec																Reserved															

2. In Volume 2, Part I, Section 11.10.3, PAL_MC_ERROR_INJECT procedure, Table 11-87, change bits 63:48 from Reserved to implementation specific as shown:

Field	Bits	Description
Reserved	47:16	Reserved
impl_spec	63:48	Processor specific error injection capabilities. Please refer to processor specific documentation for additional details.

9. MOV-to-PSR.L Reserved Field Checking

This change relaxes the architectural requirement for checking the reserved upper 32 bits on MOV-to-PSR.L, making this check implementation-specific.

Volume 3, Part I, Section 2.2, “Instruction Descriptions, “mov - Move Processor Status Register” page. In the Description, change the third paragraph from:

For move to processor status register, GR r_2 is read, bits {31:0} copied into PSR{31:0} and bits{45:32} are ignored. All bits of GR r_2 corresponding to reserved fields of the PSR must be 0 or a Reserved Register/Field fault will result.

to:

For move to processor status register, GR r_2 is read, bits {31:0} copied into PSR{31:0} and bits {63:32} are ignored. Bits {31:0} of GR r_2 corresponding to reserved fields of the PSR must be 0 or a Reserved Register/Field fault will result. An implementation may also raise Reserved Register/Field fault if bits {63:32} in GR r_2 corresponding to reserved fields of the PSR are non-zero.

10. Virtual Machine Disable

This change defines a mechanism to disable processor virtualization features.

1. 2. In Volume 2, Part I, Section 11.10.3, “PAL Procedure Specifications”, PAL_PROC_GET_FEATURES page, Table 11-103, change bit 40 to the following:



Bit	Class	Control	Description
40	Opt.	Opt	<p>Virtual Machine features implemented and enabled. When 1, PSR.vm is implemented and virtual machines features are not disabled. When 0 (<i>features_status</i>) and when the corresponding <i>features_avail</i> bit is 1, virtual machines features are implemented but are disabled. When both the <i>features_avail</i> and <i>features_status</i> bits are 0, virtual machine features are not implemented.</p> <p>If implemented and controllable, virtual machine features may be disabled by writing this bit to 0 with PAL_PROC_SET_FEATURES. However, virtual machine features cannot be re-enabled except via a reset; hence, if virtual machine features are disabled, this bit reads as 0 for both <i>features_status</i> and <i>features_control</i> (but still 1 for <i>features_avail</i>).</p>

2. Volume 2, Part I, Section 3.4, "Processor Virtualization". Add the following paragraph before the last paragraph:

Processors which support processor virtualization may provide an implementation-dependent mechanism to disable virtual machine features, see PAL_PROC_GET_FEATURES on page 2:429 for details.
3. Volume 3, Part I, Section 2.2, "Instruction Descriptions", vmsw page. In the last sentence of the last paragraph:

See Section 3.4, "Processor Virtualization" on page 2:40 and PAL_PROC_GET_FEATURES on page 2:433 for details.

add a reference to PAL_PROC_SET_FEATURES:

See Section 3.4, "Processor Virtualization" on page 2:40, PAL_PROC_GET_FEATURES and PAL_PROC_SET_FEATURES on page 2:429 and page 2:433 for details.
4. Volume 2, Part I, Section 11.8, "PAL Glossary", add the following two definitions:

Power-on

The reset event that occurs when the power input to the processor is applied and the reset input to the processor is asserted.

Reset

The reset event that occurs when the reset input to the processor is asserted.

5. Volume 2, Part I, Section 11.2.1, "PALE_RESET". In the first sentence, change the following text:

Upon receipt of a power-on reset event the processor begins executing code from...

to:

Upon receipt of a power-on/reset event the processor begins executing code from...
6. Volume 3, Part I, Chapter 3, Table 3-1, "Pseudo-code Functions." Change the description of the function "implemented_vm" from:

Returns TRUE if the processor implements the PSR.vm bit.

to:

Returns TRUE if the processor implements the PSR.vm bit (regardless of whether virtual machine features are enabled or disabled).
7. Volume 3, Part I, Chapter 3, Table 3-1 "Pseudo-code Functions", rename the function "vm_disabled" to "vmsw_disabled".
8. Volume 3, Part I, Chapter 3, Table 3-1 "Pseudo-code Function", add a new function "vm_disabled":



Function	Operation
vm_disabled	Returns TRUE if the processor implements the PSR.vm bit and virtual machine features are disabled. See Section 3.4, "Processor Virtualization" on page 2:40 in SDM and "PAL_PROC_GET_FEATURES - Get Processor Dependent Features (17)" on page 2:433 for details.

9. Volume 3, Part I, Chapter 2, VMSW I-page, change the line:

```
if (!(PSR.it == 1 && itlb_ar() == 7) || vm_disabled())
```

to:

```
if (!(PSR.it == 1 && itlb_ar() == 7) || vm_disabled() ||  
vmsw_disabled())
```

10. Volume 3, Part I, Section 2.2, "Instruction Descriptions", vmsw page. In the description section, change the last paragraph from:

Implementation of PSR.vm is optional. If it is not implemented, this instruction takes Illegal Operation fault. If it is implemented but is disabled, this instruction takes Virtualization fault when executed at the most privileged level. See Section 3.4, "Processor Virtualization" on page 2:40 and PAL_PROC_GET_FEATURES on page 2:433 for details.

to:

Implementation of PSR.vm is optional. If it is not implemented, this instruction takes Illegal Operation fault. If it is implemented but either virtual machine features or the vmsw instruction are disabled, this instruction takes Virtualization fault when executed at the most privileged level. See Section 3.4, "Processor Virtualization" on page 2:40 and PAL_PROC_GET_FEATURES on page 2:433 for details.

11. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications". In the PAL_VP_ENV_INFO page, add the following paragraph to the end of the section (just before Table 11-109):

This procedure returns unimplemented procedure when virtual machine features are disabled. See Section 3.4, "Processor Virtualization" on page 2:40 and PAL_PROC_GET_FEATURES on page 2:433 for details.

12. Volume 2, Part I, Chapter 11, Section 11.10.3, "PAL Procedure Specifications", PAL_VP_INIT_ENV page, add the following paragraph at the end of the Description:

This procedure returns unimplemented procedure when virtual machine features are disabled. See Section 3.4, "Processor Virtualization" on page 2:40 and PAL_PROC_GET_FEATURES on page 2:433 for details.

13. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_PROC_GET_FEATURES procedure, Table 11-103, add the following sentence to the end of the description for bit 54, "Enable the use of the vmsw instruction":

This bit has no effect if virtual machine features are disabled (see bit 40).

14. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_PROC_GET_FEATURES procedure, add a sentence to the following existing paragraph on the PAL_PROC_GET_FEATURES page:

For each valid *feature_set*, this procedure returns which processor features are implemented in the *features_avail* return argument, the current feature setting is in *feature_status* return argument, and the feature controllability in the *feature_control* return argument. Only the processor features which are implemented and controllable can be changed via PAL_PROC_SET_FEATURES.

to:

For each valid *feature_set*, this procedure returns which processor features are implemented in the *features_avail* return argument, the current feature setting



is in *feature_status* return argument, and the feature controllability in the *feature_control* return argument. Only the processor features which are implemented and controllable can be changed via `PAL_PROC_SET_FEATURES`. Features for which *features_avail* are 0 (unimplemented features) also have *features_status* and *features_control* of 0.

11. Removal of *pal_proc_vector* Argument from `PAL_VP_SAVE` and `PAL_VP_RESTORE`

This change simplifies `PAL_VP_SAVE` and `PAL_VP_RESTORE` implementations by removing the *pal_proc_vector* argument from these calls.

1. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", `PAL_VP_RESTORE` page:
 - a. In the Argument section, change "pal_proc_vector" and corresponding description to "Reserved" and "0".
 - b. In the Description section, remove the paragraph:

The *pal_proc_vector* parameter for `PAL_VP_RESTORE` allows the VMM to control the PAL procedure implementation-specific state to be saved. Table 11-111 shows the format of *pal_proc_vector*. When a bit is set to 1 in the vector, the implementation-specific state for the corresponding PAL procedures will be restored by `PAL_VP_RESTORE`. When a bit is set to 0 in the vector, no implementation-specific state will be restored for the corresponding PAL procedures.
 - c. Remove Table 11-111, "Format of *pal_proc_vector*".
2. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", `PAL_VP_SAVE` page:
 - a. In the Argument section, change "pal_proc_vector" and corresponding description to "Reserved" and "0".
 - b. In the Description section, remove the following paragraph:

The *pal_proc_vector* parameter for `PAL_VP_SAVE` allows the VMM to control the PAL procedure implementation-specific state to be saved. Table 11-111 on page 2:463 shows the format of *pal_proc_vector*. When a bit is set to 1 in the vector, the implementation-specific state for the corresponding PAL procedures will be saved by `PAL_VP_SAVE`. When a bit is set to 0 in the vector, no implementation-specific state will be saved for the corresponding PAL procedures.


12. Variable Frequency Mode Additions to ACPI P-states

1. In Volume 2, Part I, Section 11.6.1, "Power/Performance States (P-states)", change the following text before Figure 11-22:

`PAL_GET_PSTATE`: This procedure returns the performance index of the logical processor, relative to the highest available P-state P0 which has an index value of 100. For example, if the value returned by the procedure is 80, it indicates that the performance of the logical processor over the last time period was 20% lower than the P0 performance capability of the logical processor. The performance index is measured over the time interval since the last `PAL_GET_PSTATE` call.

to:

`PAL_GET_PSTATE`: This procedure returns the performance index of the logical processor, relative to the highest available P-state P0. A value of 100 in P0 represents the minimum processor performance in the P0 state. For example, if the value returned by the procedure is 80, this indicates that the performance



of the logical processor over the last time period was 20% lower than the minimum P0 performance. For processors that support variable P-states, it is possible for a processor to report a number greater than 100, representing that the processor is running at a performance level greater than the minimum P0 performance. For example, if the value returned by the processor is 120, it indicates that the performance of the logical processor over the last time period was 20% higher than the minimum P0 performance. The performance index is measured over the time interval since the last PAL_GET_PSTATE call with a type operand of 1. If the processor supports variable P-state performance then the PAL_PROC_SET_FEATURE procedure can be used to enable or disable this feature.

2. Add the following text to Volume 2, Part I, Section 11.6.1, "Power/Performance States (P-states)", just before section 11.6.1.1:

Some processors may support variable P-state performance where the frequency within a given P-state may vary to achieve the maximum performance for that P-state's power budget. The PAL_PROC_GET_FEATURES procedure on page 2:429 indicates if the processor supports variable P-state performance.

The performance index calculation is slightly different when a processor supports variable P-state performance, since the frequency within a P-state can vary. These frequencies for a given P-state are represented by an index value $F_{x,y}$. The value x is the P-state number and y represents a frequency point in the range from 0 to N . A value of 0 represents the minimum frequency index value for the given P-state. For example:

$F_{0,0}$ to $F_{0,N}$ - Frequency index values for the P0 state

$F_{1,0}$ to $F_{1,N}$ - Frequency index values for the P1 state

etc..

$F_{0,0}$ is the minimum frequency index for the P0 state and its value is 100. $F_{0,1}$ represents a higher frequency point for P0 and will have a value greater than 100. For example if $F_{0,1}$ frequency is 5% greater than $F_{0,0}$ it would have a value of 105.

The *performance_index* equation for P0 is calculated as follows:

$$((F_{0,0} * \text{time spent in } F_{0,0}) + (F_{0,1} * \text{time spent in } F_{0,1}) + \dots (F_{0,N} * \text{time spent in } F_{0,N})) / (\text{Total Time spent in P0})$$

For example let's say the minimum frequency of P0 is 1GHz and the maximum frequency of P0 is 1.5GHz. If we are at 1GHz for a time period of 4, 1.25GHz for a time period of 16 and 1.5GHz for a time period of 20, the average performance index is: $((100*4) + (125*16) + (150*20)) / (5+15+20) = 135$

The *performance_index* equation for other P-states can be calculated in a similar manner using their respective frequency index values.

The total *performance_index* equation for a processor with four P-states (P0, P1, P2, P3) would be:

$$\begin{aligned} & ((F_{0,0} * \text{time spent in } F_{0,0}) + (F_{0,1} * \text{time spent in } F_{0,1}) + \dots (F_{0,N} * \text{time spent in } F_{0,N})) + \\ & (F_{1,0} * \text{time spent in } F_{1,0}) + (F_{1,1} * \text{time spent in } F_{1,1}) + \dots (F_{1,N} * \text{time spent in } F_{1,N}) + \\ & (F_{2,0} * \text{time spent in } F_{2,0}) + (F_{2,1} * \text{time spent in } F_{2,1}) + \dots (F_{2,N} * \text{time spent in } F_{2,N}) + \\ & (F_{3,0} * \text{time spent in } F_{3,0}) + (F_{3,1} * \text{time spent in } F_{3,1}) + \dots (F_{3,N} * \text{time spent in } F_{3,N}) \\ & / (\text{Total Time}) \end{aligned}$$



3. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_GET_PSTATES page.

a. Change the *performance_index* return value definition from:

Unsigned integer denoting the processor performance for the time duration since the last PAL_GET_PSTATE procedure call was made. The value returned is between 0 and 100, and is relative to the performance index of the highest available P-state.

to:

Unsigned integer denoting the processor performance for the time duration since the last PAL_GET_PSTATE procedure call was made. The value returned is relative to the performance index of the highest available P-state.

b. Change the first paragraph of the Description section:

This procedure returns the performance index of the processor over the time period between the previous and the current invocations of PAL_GET_PSTATE, and is relative to the highest available P-state. For processors that belong to a software-coordinated dependency domain or a hardware-independent dependency domain, the *performance_index* value returned will correspond to the target P-state requested by the most recent PAL_SET_PSTATE procedure call.

to:

This procedure returns the performance index of the processor over the time period between the previous and the current invocations of PAL_GET_PSTATE, and is relative to the highest available P-state, P0. A value of 100 represents the minimum processor performance in the P0 state. For processors that support variable P-state performance, it is possible for a processor to report a number greater than 100, representing that the processor is running at a performance level greater than the minimum P0 performance. The PAL procedure PAL_PROC_GET_FEATURES on page 2:429 indicates if the processor supports variable P-state performance.

For processors that belong to a software-coordinated dependency domain or a hardware-independent dependency domain, the *performance_index* value returned will correspond to the target P-state requested by the most recent PAL_SET_PSTATE procedure call in cases where variable P-state performance is not supported. When variable P-states performance is supported, the *performance_index* may be higher than the target P-state requested. Please see Section 11.6.1 for more information about variable P-state performance.

c. In the Description section, change the second paragraph after Table 11-73.

If there was a thermal-throttling event or any hardware-initiated event, which affected the processor power/performance for the current time period and the accuracy of the *performance_index* value has been impacted by the event, then the procedure will return with status=1. The *performance_index* returned in this case will still have a value between 0 and 100.

to:

If there was a thermal-throttling event or any hardware-initiated event which affected the processor power/performance for the current time period and the accuracy of the *performance_index* value has been impacted by the event, then the procedure will return with status=1. The *performance_index* returned in this case will still have a value that falls within the range of




possible *performance_index* values for this processor implementation. (i.e. 0 up to the highest variable p-state *performance_index* value)

4. In Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", PAL_PROC_GET_FEATURES page, modify the bit 39 description as shown:

Bit	Class	Control	Description
39	Opt.	Req.	Variable P-state performance: A value of 1 indicates that the processor is optimizing performance for the given P-state power budget by dynamically varying the frequency, such that maximum performance is achieved for the power budget. A value of 0 indicates that P-states have no frequency variation or very small frequency variations for their given power budget.

13. PAL_MC_DYNAMIC_STATE Changes

1. Volume 2, Part I, Section 11.10.1, "PAL Procedure Summary", Table 11-38. Change the PAL_MC_DYNAMIC_STATE "Mode" value from "Phys." to "Both".
2. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications". Update the PAL_MC_DYNAMIC page as follows:



PAL_MC_DYNAMIC_STATE – Returns Dynamic Processor State (24)

Purpose: Returns the Machine Check Dynamic Processor State.

Calling Conv: Static Registers Only

Mode: Physical and Virtual

Buffer: Not dependent

Arguments:	Arguments	Description
	index	Index of PAL_MC_DYNAMIC_STATE within the list of PAL procedures.
	info_type	Unsigned 64-bit value indicating the type of information to return
	dy_buffer	64-bit pointer to a buffer aligned on an 8-byte boundary
Returns:	Reserved	0
	Return Value	Description
	status	Return status of the PAL_MC_DYNAMIC_STATE procedure.
	max_size	Maximum size in bytes of the data that can be returned by this procedure for this processor family.
	Reserved	0
Status:	Reserved	0
	Status Value	Description
	0	Call completed without error
	-1	Unimplemented procedure
	-2	Invalid argument
	-3	Call completed with error

Description: The *info_type* input argument designates the type of information the procedure will return. When *info_type* is 0, the procedure returns the maximum size in bytes of processor dynamic state that can be returned for this processor family in the *max_size* return value.

When *info_type* is 1, the procedure will copy processor dynamic state into memory pointed to by the input argument *dy_buffer*. This copy will occur using the addressing attributes used to make the procedure call (physical or virtual) and the caller needs to ensure the *dy_buffer* input pointer matches this addressing attribute.

The amount of data returned can vary depending on the state of the machine when called and may not always return the maximum size for every call. The amount of data returned is provided in the processor state parameter field *dsize*. Please see Table 11-7 for more information on the processor state parameter. The caller of the procedure needs to ensure that the buffer is large enough to handle the *max_size* that is returned by this procedure.

The contents of the processor dynamic state is implementation dependent. Portions of this information may be cleared by the PAL_MC_CLEAR_LOG procedure. This procedure should be invoked before PAL_MC_CLEAR_LOG to ensure all the data is captured.


14. Min-State Save Area Size Change

1. In Volume 2, Part I, Section 11.2.2.1, "Definition of SALE_ENTRY State Parameter", change the reset hand-off state to return a min-state save area size.

Page 2:282 - Modify the existing SALE_ENTRY state parameter for reset and recovery check to pass the size of the min-state save area using some previously reserved fields

- a. Modify Figure 11-8 to add a new field called *min-state_size* as shown:

Figure 11-8. SALE_ENTRY State Parameter



- b. Add a bullet at the end of Section 11.2.2.1, after the sentence "For the case of RECOVERY CHECK, authentication of PAL_A and PAL_B should be completed before call to SAL_ENTRY":
 - *min-state_size* - An 8-bit field indicating the size in kilobytes (KB) of the min-state save area required for this implementation. A value of zero indicates a size of 4 KB. A value greater than zero indicates the actual size in KB of the min-state save area required for this implementation. Values of 1-4 are reserved. For more information about the min-state save area, please refer to Section 11.3.2.3.
2. Volume 2, Part I, Section 11.3.2.3, "Processor Min-state Save Area Layout".
 - a. Change the first paragraph of Section 11.3.2.3 from:

The processor min-state save area is 4 KB in size and must be in an uncacheable region. The first 1 KB of this area is architectural state needed by the PAL code to resume during MCA and INIT events (architected min-state save area + reserved). The remaining 3KB is a scratch buffer reserved exclusively for PAL use, therefore SAL and OS must not use this area. The layout of the processor min-state save area is shown in Figure 11-13.

to:

The processor min-state save area is minimally 4 KB in size, but an implementation may require larger sizes. The reset hand-off state indicates if a size greater than 4 KB is required and also provides the required size. Please refer to Section 11.2.2.1 for more information on the reset hand-off state. The required size is referred to as MIN_STATE_REQ. The min-state save area is required to be in an uncacheable region. The first 1 KB of this area is architectural state needed by the PAL code to resume during MCA and INIT events (architected min-state save area + reserved). The remaining space in the buffer is a scratch space reserved exclusively for PAL use, therefore SAL and OS must not use this area. The layout of the processor min-state save area is shown in Figure 11-13.
 - b. Figure 11-13 needs to be modified in two places to use the MIN_STATE_REQ variable:

Change "Min-state save ptr + 4KB" to "Min-state save ptr + MIN_STATE_REQ" and change "3KB" to "MIN_STATE_REQ - 1KB".



- c. Change the third paragraph of Section 11.3.2.3 from:
The base address of the min-state save area must minimally be aligned on a 512-byte boundary, but larger alignments like 4 KB are fine.
to:
The base address of the min-state save area must minimally be aligned on a 512-byte boundary with, but larger alignments are allowed.
- 3. Change the PAL_MC_REGISTER_MEM procedure as shown:



PAL_MC_REGISTER_MEM – Register Memory with PAL for Machine Check and Init (27)

Purpose: Registers a platform dependent location with PAL to which it can save minimal processor state in the event of a machine check or initialization event.

Calling Conv: Static Registers Only

Mode: Physical

Buffer: Not dependent

Arguments:	Argument	Description
	index	Index of PAL_MC_REGISTER_MEM within the list of PAL procedures.
	address	Physical address of the buffer to be registered with PAL.
	size	Unsigned integer indicating the size in kilobytes (KB) of the buffer passed. This input argument is only required when passing in a size greater than 4KB. The implementation indicates when a size greater than 4KB is required at the reset hand-off. Refer to Section 11.2.2.1 for more information.
	Reserved	0
Returns:	Return Value	Description
	status	Return status of the PAL_MC_REGISTER_MEM procedure.
	req_size	Returns the required size of the min-state save area in kilobytes (KB) if the 'size' input argument did not match the required size for this implementation.
	Reserved	0
	Reserved	0
Status:	Status Value	Description
	0	Call completed without error
	-2	Invalid argument
	-3	Call completed with error

Description: PAL places the address passed in the XRO register, which is used by PAL as the min-state save area

This procedure is used to register with PAL an uncacheable min-state save area memory buffer that is used for machine check and initialization event handling. The size of the min-state save area is either 4 KB or a larger size that is indicated in the reset hand-off state described in Section 11.2.2.1. The input argument *size* indicates the size of the min-state save buffer in kilobytes (KB) when it is greater than 4 KB. If the *size* input argument does not match the required size, the procedure returns an invalid argument return status and a min-state area is not registered. The procedure will also return the required size of the min-state save area in the *reg_size* return value.

The layout of the min-state save area is defined in Section 11.3.2.3 “Processor Min-state Save Area Layout” on page 2:294. The address passed has a minimum alignment requirement of 512-bytes.

15. Data Speculation Disable

This specification change provides a mechanism for disabling data speculation to force code execution to be more reproducible.

4. Volume 2, Part I, Section 11.10.3, PAL_PROC_GET_FEATURES page. Add a new bit 35 to Table 11-103, “Processor Features” as shown:



Bit	Class	Control	Description
35	Opt.	Req	Disable data speculation and the ALAT. When 1, data speculation checks (chk.a) always fail (i.e., always branch to the target address), thus triggering recovery code; check loads (ld.c) always re-load the target register. When 0, data speculation works as normal.
34:0	N/A	N/A	Reserved

16. Interruption Instruction Bundle Registers

This specification change defines Interruption Instruction Bundle registers, which provide instruction bundle information for certain IVA-based faults and traps.

1. Volume 2, Part I, Chapter 3 “System State and Programming Model”.
 - a. In Figure 3-1, “System Register Model”, add CR26 IIB0 and CR27 IIB1 to the diagram.
 - b. In Table 3-3, “Control Registers” add CR26 and CR27 rows as shown:

	Register	Name	Description	Serialization Required
Interruption Control Registers	CR26	IIB0	Interruption Instruction Bundle 0	implied ^c
	CR27	IIB1	Interruption Instruction Bundle 1	implied ^c
Reserved	CR28-63		reserved	

- c. Section 3.3.5, “Interruption Control Registers”. Change the first sentence from:

Registers CR16 - CR25 record information at the time of an interruption...


 to:

Registers CR16 - CR27 record information at the time of an interruption...
- d. Add a new Section 3.3.5.10 after the current Section 3.3.5.9:

3.3.5.10 Interruption Instruction Bundle Registers (IIB0-1 - CR26,27)


On an interruption and if PSR.ic is 1, the IIB registers receive the 16-byte instruction bundle corresponding to the interruption. The bundle reported in the IIB registers is the bundle exactly as it was fetched for execution of the instruction which raised the interruption. Figure 3-16 shows the format of the IIB0 and IIB1 registers. For details on instruction bundle format, see Vol. 1, Section 3.3, “Instruction Encoding Overview”.

Figure 3-16. Interruption Instruction Bundle Registers (IIB0-1, – CR26, 27)



If the interruption is a fault, the IIB registers record the instruction bundle pointed to by IIP. If the interruption is a trap, the IIB registers record the instruction bundle pointed to by IIPA.

The IIB registers only provide valid interruption bundle information on certain IVA-based faults and traps. Please refer to Table 8-1, “Writing of Interruption Resources by Vector” and corresponding interruption vector pages in Section 8.3, “Interruption



Vector Definition” for information on which faults and traps these registers are valid. For faults and traps that indicate IIB is not valid, updates to the register may occur, but the information is undefined.

For IA-32 interruptions, instruction bundle information is not provided and the values in IIB registers are undefined.

The IIB registers are not supported on all processor implementations. Software can call PAL_PROC_GET_FEATURES to determine the availability of this feature. The IIB registers are reserved when this feature is not supported.

2. Volume 2, Part I, Chapter 5, “Interruptions”.

- a. Section 5.2, last paragraph, add the following to the end of the list of control registers in parentheses:

, IIB0-1

- b. In Section 5.5, change the line after “1. (If PSR.ic is 0)” from:

IPSR, IIP, IIPA, and IFS.v are unchanged.

to:

IPSR, IIP, IIPA, IIB0-1, and IFS.v are unchanged.

- c. In Section 5.5, change the fourth bullet under “If PSR.ic is 1” from:

- The interruption resources IFA, IIM, IHA, and ITIR are written with information...

to:

- The interruption resources IFA, IIB0-1, IIM, IHA, and ITIR are written with information...

- d. In Section 5.5, change the line after “If PSR.ic is in-flight” from:

Interruption state may or may not be collected in IIP, IPSR, IIPA, ITIR, IFA, IIM and IHA.

to:

Interruption state may or may not be collected in IIP, IPSR, IIPA, ITIR, IFA, IIM, IIB0-1 and IHA.

3. Volume 2, Part I, Chapter 8, “Interruption Vector Descriptions”


- a. Replace Table 8-1 with the following:

Table 8-1. Writing of Interruption Resources by Vector


Interruption Resource	IIP, IPSR, IIPA, IFS.v		IFA		ITIR		IHA		IIM		ISR		IIB0, IIB1
PSR.ic at time of interruption	0	1	0	1	0	1	0	1	0	1	0	1	0
Alternate Data TLB vector													
Alternate Data TLB fault	N/A ^a	W ^b	N/A	W	N/A	W	N/A	x ^c	N/A	x	N/A	W	N/A
IR Alternate Data TLB fault	N/A	W	N/A	W	N/A	W	N/A	x	N/A	x	N/A	W	N/A
Alternate Instruction TLB vector													
Alternate Instruction TLB fault	- ^d	W	-	W	-	W	x	x	x	x	W	W	-
Break Instruction vector													
Break Instruction fault	-	W	x	x	x	x	x	x	-	W	W	W	-
Data Access Rights vector													
Data Access Rights fault	-	W	-	W	-	W	x	x	x	x	W	W	-
IR Data Access Rights fault	-	W	-	W	-	W	x	x	x	x	W	W	-



Interrupt Resource	IIP, IPSR, IIPA, IFS.v		IFA		ITIR		IHA		IIM		ISR		IIB0, IIB1
PSR.ic at time of interruption	0	1	0	1	0	1	0	1	0	1	0	1	0
Data Access-Bit vector													
Data Access Bit fault	-	W	-	W	-	W	x	x	x	x	W	W	-
IR Data Key Miss fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Data Key Miss vector													
Data Key Miss fault	-	W	-	W	-	W	x	x	x	x	W	W	-
IR Data Key Miss fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Data Nested TLB vector													
Data Nested TLB fault	-	N/A	-	N/A	-	N/A	-	N/A	x	N/A	-	N/A	-
IR Data Nested TLB fault	-	N/A	-	N/A	-	N/A	-	N/A	x	N/A	-	N/A	-
Data TLB vector													
Data TLB fault	N/A	W	N/A	W	N/A	W	N/A	W	N/A	x	N/A	W	N/A
IR Data TLB fault	N/A	W	N/A	W	N/A	W	N/A	W	N/A	x	N/A	W	N/A
Debug vector													
Data Debug fault	-	W	-	W	x	x	x	x	x	x	W	W	-
Instruction Debug fault	-	W	-	W	x	x	x	x	x	x	W	W	-
IR Data Debug fault	-	W	-	W	x	x	x	x	x	x	W	W	-
Dirty-Bit vector													
Data Dirty Bit fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Disabled FP-Register vector													
Disabled Floating-Point Register fault	-	W	x	x	x	x	x	x	x	x	W	W	-
External Interrupt vector													
External Interrupt	-	W	x	x	x	x	x	x	x	x	W	W	-
Floating-point Fault vector													
Floating-Point Exception fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Floating-point Trap vector													
Floating-Point Exception trap	-	W	x	x	x	x	x	x	x	x	W	W	-
General Exception vector													
Disabled ISA Transition fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Illegal Dependency fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Illegal Operation fault	-	W	x	x	x	x	x	x	x	x	W	W	-
IR Unimplemented Data Address fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Privileged Operation fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Privileged Register fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Reserved Register/Field fault	-	W	x	x	x	x	x	x	x	x	W	W	-
Unimplemented Data Address fault	-	W	x	x	x	x	x	x	x	x	W	W	-
IA-32 Exception vector	N/A	W	N/A	x	N/A	x	N/A	x	N/A	x	N/A	W	N/A
IA-32 Intercept vector	N/A	W	N/A	x	N/A	x	N/A	x	N/A	W	N/A	W	N/A
IA-32 Interrupt vector	N/A	W	N/A	x	N/A	x	N/A	x	N/A	x	N/A	W	N/A
Instruction Access Rights vector													
Instruction Access Rights fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Instruction Access-Bit vector													
Instruction Access Bit fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Instruction Key Miss vector													




Interruption Resource	IIP, IPSR, IIPA, IFS.v		IFA		ITIR		IHA		IIM		ISR		IIB0, IIB1
PSR.ic at time of interruption	0	1	0	1	0	1	0	1	0	1	0	1	0
Instruction Key Miss fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Instruction TLB vector													
Instruction TLB fault	-	W	-	W	-	W	-	W	x	x	W	W	-
Key Permission vector													
Data Key Permission fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Instruction Key Permission fault	-	W	-	W	-	W	x	x	x	x	W	W	-
IR Data Key Permission fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Lower-Privilege Transfer Trap vector													
Unimplemented Instruction Address fault	-	W	x	W	x	x	x	x	x	x	W	W	-
Lower-Privilege Transfer trap	-	W	x	x	x	x	x	x	x	x	W	W	-
Unimplemented Instruction Address trap	-	W	x	x	x	x	x	x	x	x	W	W	-
NaT Consumption vector													
Data NaT Page Consumption fault	-	W	-	W	x	x	x	x	x	x	W	W	-
Instruction NaT Page Consumption fault	-	W	-	W	x	x	x	x	x	x	W	W	-
IR Data NaT Page Consumption fault	-	W	-	W	x	x	x	x	x	x	W	W	-
Register NaT Consumption fault	-	W	-	x	x	x	x	x	x	x	W	W	-
Page Not Present vector													
Data Page Not Present fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Instruction Page Not Present fault	-	W	-	W	-	W	x	x	x	x	W	W	-
IR Data Page Not Present fault	-	W	-	W	-	W	x	x	x	x	W	W	-
Single Step Trap vector													
Single Step trap	-	W	x	x	x	x	x	x	x	x	W	W	-
Speculation vector													
Speculative Operation fault	-	W	x	x	x	x	x	x	-	W	W	W	-
Taken Branch Trap vector													
Taken Branch trap	-	W	x	x	x	x	x	x	x	x	W	W	-
Unaligned Reference vector													
Unaligned Data Reference fault	-	W	-	W	x	x	x	x	x	x	W	W	-
Unsupported Data Reference vector													
Unsupported Data Reference fault	-	W	-	W	x	x	x	x	x	x	W	W	-
VHPT Translation vector													
IR VHPT Data fault	N/A	W	N/A	W	N/A	W	N/A	W	N/A	x	N/A	W	N/A
VHPT Data fault	N/A	W	N/A	W	N/A	W	N/A	W	N/A	x	N/A	W	N/A
VHPT Instruction fault	N/A	W	N/A	W	N/A	W	N/A	W	N/A	x	N/A	W	N/A
Virtual External Interrupt vector													
Virtual External Interrupt	-	W	x	x	x	x	x	x	x	x	W	W	-



Interrupt Resource	IIP, IPSR, IIPA, IFS.v		IFA		ITIR		IHA		IIM		ISR		IIB0, IIB1
PSR.ic at time of interruption	0	1	0	1	0	1	0	1	0	1	0	1	0
Virtualization vector													
Virtualization fault	-	W	x	x	x	x	x	x	x	x	W	W	-

Notes:

- "N/A" indicates that this cannot happen.
 - "W" indicates that the resource is written with a new value.
 - "x" indicates that the resource may or may not be written; whether it is written and with what value is implementation specific.
 - "-" indicates that the resource is not written.
- Volume 2, Part I, Section 8.3. On page 2:163, "VHPT Translation vector (0x0000)" in the Parameters section, add the following after "ITIR":
IIB0, IIB1 - If implemented, for VHPT Data faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR VHPT Data and VHPT Instruction faults. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - Volume 2, Part I, Section 8.3. On page 2:165, "Instruction TLB vector (0x0400)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - Volume 2, Part I, Section 8.3. On page 2:166, "Data TLB vector (0x0800)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, for Data TLB faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data TLB faults. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - Volume 2, Part I, Section 8.3. On page 2:167, "Alternate Instruction TLB vector (0x0c00)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instructions Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - Volume 2, Part I, Section 8.3. On page 2:168, "Alternate Data TLB vector (0x1000)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, for Alternate Data TLB faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Alternate Data TLB faults. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - Volume 2, Part I, Section 8.3. On page 2:169, "Data Nested TLB vector (0x1400)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers are unchanged from their previous values. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.



- h. Volume 2, Part I, Section 8.3. On page 2:170, "Instruction Key Miss vector (0x1800)" in the Parameters section, add the following after "IFA":
 - IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- i. Volume 2, Part I, Section 8.3. On page 2:171, "Data Key Miss vector (0x1c00)" in the Parameters section, add the following after "IFA":
 - IIB0, IIB1 - If implemented, for Data Key Miss faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data Key Miss faults. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- j. Volume 2, Part I, Section 8.3. On page 2:172, "Dirty-Bit vector (0x2000)" in the Parameters section, add the following after "IFA":
 - IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- k. Volume 2, Part I, Section 8.3. On page 2:173, "Instruction Access-Bit vector (0x2400)" in the Parameters section, add the following after "IFA":
 - IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- l. Volume 2, Part I, Section 8.3. On page 2:174, "Data Access-Bit vector (0x2800)" in the Parameters section, add the following after "IFA":
 - IIB0, IIB1 - If implemented, for Data Access Bit faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data Access Bit faults. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- m. Volume 2, Part I, Section 8.3. On page 2:175, "Break Instruction vector (0x2c00)" in the Parameters section, add the following after "IIM":
 - IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- n. Volume 2, Part I, Section 8.3. On page 2:176, "External Interrupt vector (0x3000)" in the Parameters section, add the following after "IVR":
 - IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- o. Volume 2, Part I, Section 8.3. On page 2:177, "Virtual External Interrupt vector (0x3400)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
 - IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- p. Volume 2, Part I, Section 8.3. On page 2:178, "Page Not Present vector (0x5000)" in the Parameters section, add the following after "ITIR":
 - IIB0, IIB1 - If implemented, for Data Page Not Present faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers



- are undefined for IR Data Page Not Present and Instruction Page Not Present faults. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- q. Volume 2, Part I, Section 8.3. On page 2:179, "Key Permission vector (0x5100)" in the Parameters section, add the following after "ITIR":
IIB0, IIB1 - If implemented, for Data Key Permission faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data Key Permission and Instruction Key Permission faults. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - r. Volume 2, Part I, Section 8.3. On page 2:180, "Instruction Access Rights vector (0x5200)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - s. Volume 2, Part I, Section 8.3. On page 2:181, "Data Access Rights vector (0x5300)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, for Data Access Rights faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data Access Rights faults. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - t. Volume 2, Part I, Section 8.3. On page 2:182, "General Exception vector (0x5400)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP for the following faults:
 - Illegal Operation fault
 - Illegal Dependency fault
 - Privileged Operation fault
 - Disabled Instruction Set Transition fault
 - Reserved Register/Field fault
 - Unimplemented Data Address fault
 - Privileged Register faultThe IIB registers are undefined for IR Unimplemented Data Address faults. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - u. Volume 2, Part I, Section 8.3. On page 2:184, "Disabled FP-Register vector (0x5500)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interrupton Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
 - v. Volume 2, Part I, Section 8.3. On page 2:185, "NaT Consumption vector (0x5600)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
IIB0, IIB1 - If implemented, for Register NaT Consumption and Data NaT Page Consumption faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data NaT Page




Consumption and Instruction NaT Page Consumption faults. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.

- w. Volume 2, Part I, Section 8.3. On page 2:187, "Speculation vector (0x5700)" in the Parameters section, add the following after "IIM":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- x. Volume 2, Part I, Section 8.3. On page 2:189, "Debug vector (0x5900)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
IIB0, IIB1 - If implemented, for Data Debug faults, the IIB registers contain the instruction bundle pointed to by IIP. The IIB registers are undefined for IR Data Debug and Instruction Debug faults. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- y. Volume 2, Part I, Section 8.3. On page 2:190, "Unaligned Reference vector (0x5a00)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- z. Volume 2, Part I, Section 8.3. On page 2:191, "Unsupported Data Reference vector (0x5b00)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- aa. Volume 2, Part I, Section 8.3. On page 2:192, "Floating-point Fault vector (0x5c00)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ab. Volume 2, Part I, Section 8.3. On page 2:193, "Floating-point Trap vector (0x5d00)" in the Parameters section, add the following after "IFA":
IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIPA. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ac. Volume 2, Part I, Section 8.3. On page 2:194, "Lower-Privilege Transfer Trap vector (0x5e00)" in the Parameters section, add the following after "Note:"
IIB0, IIB1 - If implemented, for Lower-Privilege Transfer traps, the IIB registers contain the instruction bundle pointed to by IIPA. The IIB registers are undefined for Unimplemented Instruction Address faults. Please refer to Section 3.3.5.10, "Interruption Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ad. Volume 2, Part I, Section 8.3. On page 2:196, "Taken Branch Trap vector (0x5f00)" in the Parameters section, add the following after "Note:"



- IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIPA. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ae. Volume 2, Part I, Section 8.3. On page 2:197, "Single Step Trap vector (0x6000)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
- IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIPA. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- af. Volume 2, Part I, Section 8.3. On page 2:198, "Virtualization vector (0x6100)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
- IIB0, IIB1 - If implemented, the IIB registers contain the instruction bundle pointed to by IIP. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ag. Volume 2, Part I, Section 8.3. On page 2:199, "IA-32 Exception vector (0x6900)" in the Parameters section, add the following after "IFA":
- IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ah. Volume 2, Part I, Section 8.3. On page 2:200, "IA-32 Intercept vector (0x6a00)" in the Parameters section, add the following after "IIM":
- IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
- ai. Volume 2, Part I, Section 8.3. On page 2:201, "IA-32 Interrupt vector (0x6b00)" in the Parameters section, add the following after "IIP, IPSR, IIPA, IFS":
- IIB0, IIB1 - If implemented, the IIB registers are undefined. Please refer to Section 3.3.5.10, "Interrupt Instruction Bundle Registers (IIB0-1 - CR26, 27)" on page 2:39 for details on the IIB registers.
4. Volume 2, Part I, Chapter 10, "Itanium® Architecture-based Operating System Interaction Model with IA-32 Applications".
- a. Table 10-1, "IA-32 System Register Mapping: Under "Control Registers", change the following:
- IFA, IIP, IPSR, ISR, IIM, IIPA, ITTR, IHA, IFS, IVA
- to:
- IFA, IIP, IPSR, ISR, IIM, IIPA, ITIR, IHA, IIB0-1, IFS, IVA
5. Volume 2, Part I, Chapter 11 "Processor Abstraction Layer".
- a. Table 11-14, "Virtual Processor Descriptor (VPD)". In footnote (g), add ", viib0-1" to the end of the list of control registers in parentheses
- b. Section 11.7.4.1.2, "Interrupt Control Register Read Optimization". In the first paragraph, change:
- (vipr, visr, viip, vifa, vitir, viipa, vifs, viim, viha)
- to:
- (vipr, visr, viip, vifa, vitir, viipa, vifs, viim, viha, viib0-1)
- c. Section 11.7.4.1.2, "Interrupt Control Register Read Optimization".



In Table 11-23, “Synchronization Requirements for Interruption Control Register Read Optimization”, add “, viib0-1” to the end of the list.

d. Section 11.7.4.1.3, “Interruption Control Register Write Optimization”.

In the first paragraph, change:

(vipsr, visr, viip, vifa, vitir, viipa, vifs, viim, viha)

to:

(vipsr, visr, viip, vifa, vitir, viipa, vifs, viim, viha, viib0-1)

e. Section 11.7.4.1.3, “Interruption Control Register Write Optimization”.

In Table 11-23, “Synchronization Requirements for Interruption Control Register Read Optimization”, add “, viib0-1:” to the end of the list.

f. Section 11.10.2.2.2, “System Registers”. In Table 11-45, “System Register Conventions”, add the following row under “IHA”:

Name	Description	Class
IIB0-1	Interruption Instruction Bundle Registers	scratch

g. Section 11.10.3, PAL_ENTER_IA_32_ENV procedure. In Table 11-71, “Register Values at IA-32 System Environment Termination, in the following row:

Intel® Itanium® Register	IA-32 State	Description
IFA, IIP, IPSR, ISR, IIM, IIPA, ITTR, IHA, IFS, IVA, GPTA, ITM, IVR, TPR, IRR0-3, ITV, PMV, LRR0, LRR1, CMCV		Undefined

add “IIB0-1,” after “IHA” as shown:

Intel® Itanium® Register	IA-32 State	Description
IFA, IIP, IPSR, ISR, IIM, IIPA, ITTR, IHA, IIB0-1, IFS, IVA, GPTA, ITM, IVR, TPR, IRR0-3, ITV, PMV, LRR0, LRR1, CMCV		Undefined

h. Section 11.10.3, PAL_PROC_GET_FEATURES procedure, Modify bit 35 of Table 11-103, “Processor Features” as shown:

Bit	Class	Control	Description
35	Opt.	No	Interruption Instruction Bundle interruption registers (IIB0, IIB1) implemented. Denotes whether IIB registers are implemented. This feature may only be interrogated by PAL_PROC_GET_FEATURES. It may not be enabled or disabled by PAL_PROC_SET_FEATURES. The corresponding argument is ignored.
34:0	N/A	N/A	Reserved

i. Section 11.7.1.1, “Virtualization Controls.” In Table 11-15, “Virtualization Acceleration Control (vac) Fields”, change the two instances of “(CR16-25)” to “(CR16-27)”.



- j. Section 11.7.2, "Interruption Handling in a Virtual Environment". Three paragraphs after Table 11-17, change the following bullet items from:
- mov-from-interruption-CR (CRs 16, 17, 19-25)
 - mov-to-interruption-CR (CRs 16, 17, 19-25)
- to:
- mov-from-interruption-CR (CRs 16, 17, 19-27)
 - mov-to-interruption-CR (CRs 16, 17, 19-27)
6. Volume 2, Part II, Chapter 3, "Interruptions and Serialization".
- a. Section 3.3.2, "Interruption Register State". Add the following to the list of control registers after "IFS":
- IIB0, IIB1 - Contain the 16-byte instruction bundle related to the interruption. Note that the IIB registers do not provide bundle information for all interruptions and are not supported on all processor implementations; please refer to Chapter 8, "Interruption Vector Descriptions" for details. Software can use the instruction bundle information for debug and emulation purposes.
- b. Section 3.4.2, "Heavyweight Interruptions". In numbered lists #1, #4, and #17 add ", IIB0-1" to the list of control registers in parentheses.
7. Volume 3, Part I, Chapter 5, "Resource and Dependency Semantics".
- a. Section 5.4, "Support Tables". In Table 5-5, "Instruction Classes". Add the following row after "mov-from-CR-IHA":

Class	Events/Instructions
mov-from-CR-IIB	mov-from-CR[Field(cr3) in {IIB0 IIB1}]

- b. Section 5.4, "Support Tables". In Table 5-5, "Instruction Classes". Add the following row after "mov-to-CR-IHA":

Class	Events/Instructions
mov-to-CR-IIB	mov-to-CR[Field(cr3) in {IIB0 IIB1}]

- c. Section 5.3.2, "RAW Dependency Table". In Table 5-2, "RAW Dependencies Organized by Resource" add the following row after "CR[IHA]":

Resource Name	Writers	Readers	Semantics of Dependency
CR[IIB%], % in 0 - 1	mov-to-CR-IIB	mov-from-CR-IIB	data

- d. Table 5-3, "WAW Dependencies Organized by Resource". Add the following row after "CR[IHA]":


Resource Name	Writers	Semantics of Dependency
CR[IIB%], % in 0 - 1	mov-to-CR-IIB	impliedF

8. Volume 3, Part I, Section 2.2, "Instruction Descriptions", "mov cr" instruction page. In the last paragraph of the Description section, change "(CR16-CR25)" to "(CR16-CR27)"

17. Data-Poisoning Promotion Changes

This specification changes the Processor State Parameter Hand-off State when data-poisoning promotion is enabled by setting PAL_PROC_SET_FEATURES bit 53.

1. Volume 2, Part I, Section 11.10.3, PAL_PROC_GET_FEATURES procedure. In Table 11-103, change the description of bit 53 from:



Enable MCA signaling on data-poisoning event detection. When 0, a CMCI will be signaled on error detection. When 1, an MCA will be signaled on error detection. If this feature is not supported, then the corresponding argument is ignored when calling PAL_PROC_SET_FEATURES. Note that the functionality of this bit is independent of the setting in bit 60 (Enable CMCI promotion), and that the bit 60 setting does not affect CMCI signaling for data-poisoning related events.

to:

Enable MCA signaling on unconsumed data-poisoning event detection. When 0, a CMCI will be signaled on error detection. When 1, an MCA will be signaled on error detection. Note that the reported error severity depends on which method is chosen for signaling; see section 11.3.2.3 for details. If this feature is not supported, then the corresponding argument is ignored when calling PAL_PROC_SET_FEATURES. Note that the functionality of this bit is independent of the setting in bit 60 (Enable CMCI promotion), and that the bit 60 setting does not affect CMCI signaling for data-poisoning related events.

2. Add a new Section 11.3.2.3, "Multiprocessor Rendezvous Requirements for Handling Machine Checks" just after current Section 11.3.2.2:

11.3.2.3 Unconsumed Data-Poisoning Event Handling

The transfer/access of information between levels of the cache/memory hierarchy where the data has an uncorrectable error and is therefore marked as poison may raise error reporting events. If the processor being reported to is not a consumer of the data in question, then the error is termed a "unconsumed data-poisoning event".

Unconsumed data-poisoning events are by default reported as a CMC and can optionally be promoted to an MCA via bit 53 of feature_set 0 of PAL_PROC_SET_FEATURES. When they are signaled as a CMC the PSP.cm is set to 1 to indicate that the error has been corrected (in the sense that the line has been marked with poison, preventing any silent data corruption).


If bit 53 is 1, unconsumed data-poisoning events are reported as MCAs. The caller can set bit 53 to 1 in order to handle unconsumed data-poisoning events immediately as uncorrected errors (in the sense that the data in question has been lost). PSP settings for a data-poisoning event with bit 53 equal to 1 are given in the table below. See also Table 11-8.

PSP bit settings for unconsumed data-poisoning events on MCA:

cm	us	ci	co	sy
0	0	1	1	0

When promotion is enabled (bit 51 is 1), and a continuable data-poisoning event is indicated (i.e., the PSP bits are set as in the above table, and either or both of cache_check.dp or bus_check.dp are 1), and if no other MCAs occur at the same time (i.e., no other errors are indicated in the error information from PAL_MC_ERROR_INFO), the interrupted process is always continuable. Promotion to MCA with bit 53 allows the OS to take proactive measures to recover from the poisoned data, but this is not required in order for the interrupted process to be continuable.

3. In Section 11.10.3, PAL_MC_ERROR_INFO description, Table 11-82, "cache_check fields", change the description of the dp field (bit 23) from:
A multiple-bit error was detected, and data was poisoned for the corresponding cache line during castout.



to:

An uncorrectable (typically multiple-bit) error was detected, and data was poisoned for the corresponding cache line, without any corrupted data being consumed (i.e., no corrupted data has been copied to processor registers).


18. ACPI P-State Clarifications

1. Volume 2, Part I, Section 11.6.1:

- a. In the first paragraph, replace:
(hence to be referred as P-states)
with:
(hence to be referred to as P-states)
- a. In the second paragraph, replace the sentence:
Successive P-states continue to have reduced performance capabilities and reduced power consumption than the corresponding lower state.
with:
Successive P-states continue to have reduced performance capabilities and reduced power consumption.
- b. Just after Figure 11-21, "Example of a P-state Transition Policy" and before the paragraph that begins, "The concept of P-states...", add this section heading:

11.6.1.1 Power Dependency Domains

- c. In the paragraph that begins, "The concept of P-states...", replace the last sentence:
To allow the architecture definition to comprehend for multi-threaded/multi-core designs, we define the concept of dependency domain and coordination mechanisms.
with:
To allow the architecture definition to comprehend multi-threaded/multi-core designs, we define the concept of dependency domain and coordination mechanisms.
- d. In the paragraph that begins "A dependency domain is comprised...", replace the sentence:
As an example, a processor package comprising of two cores controlled by the same clock and power distribution network are part of the same dependency domain, since changing either the operating frequency or voltage will affect power consumption and performance for both cores.
with:
As an example, a processor package comprised of two cores controlled by the same clock and power distribution network are part of the same dependency domain, since changing either the operating frequency or voltage will affect power consumption and performance for both cores.
- e. In the paragraph that begins, "A dependency domain is comprised...", replace the sentence:
Software can utilize P-states to affect changes in the domain parameters.
with:
Software can utilize P-states to effect changes in the domain parameters.
- f. In the paragraph that begins, "A software-coordinated dependency domain...", replace the beginning of the sentence:
A software-coordinated dependency domain relies on the software...



with:

A software-coordinated dependency domain (SCDD) relies on the software...

- g. In the paragraph that begins, "A software-coordinated dependency domain...", replace the sentence:

As an example, let us assume that the software-coordinated dependency domain consisted of two cores with the same clock and power distribution networks and the intent of the software policy was to lower power/performance only when the workload utilization was low on both cores.

with:

As an example, let us assume that the SCDD consisted of two cores with the same clock and power distribution networks and the intent of the software policy was to lower power/performance only when the workload utilization was low on both cores.

- h. In the paragraph that begins, "A software-coordinated dependency domain...", replace the sentence:

This transition would simultaneously reduce performance and power dissipation for both cores, and would result in both cores operating at the same P-state.

with:

This transition would simultaneously reduce performance and power dissipation for both cores, and would result in both cores operating at the same lower P-state.

- i. In the paragraph that begins, "A hardware-coordinated dependency domain...", replace the beginning of the sentence:

A hardware-coordinated dependency domain relies on the software...

with:

A hardware-coordinated dependency domain (HCDD) relies on the software...

- j. In the paragraph that begins, "A hardware-coordinated dependency domain...", replace the sentence:

Software can make independent P-state change requests on individual processors, recognizing that hardware is responsible for the required coordination with other processors in the same hardware-coordinated dependency domain.

with:

Software can make independent P-state change requests on individual processors, recognizing that hardware is responsible for the required coordination with other processors in the same HCDD.

- k. In the paragraph that begins, "A hardware-coordinated dependency domain...", replace the following text:

As an example, let us assume that the hardware-coordinated dependency domain consisted of two cores with the same clock and power distribution networks, and that there were also some other techniques to affect power and performance which were local to each logical processor. When software initiates a P-state transition on the first core, hardware would use only the local parameters to carry out the request. When software requests the same P-state change on the second core, then hardware can undo the changes to the local parameters for the first core, and then initiate changes to the domain parameters, which would allow both cores to operate at the same P-state.



with:

Domain parameters are set by hardware according to the highest requested power/performance level (that is, the lowest numbered P-state) of the logical processors in the power domain. As an example, let us assume that the HCDD consisted of two cores with the same clock and power distribution networks, and that there were also some other techniques to affect power and performance which were local to each logical processor. Let us also assume that software has initially set both cores to the P0 state. When software initiates a P-state transition to P1 (which is a lower power/performance level) on the first core, hardware would use only the local parameters to carry out the request, and the domain parameters would remain at P0. Suppose software on the second core then initiates a P-state transition to P3. Hardware would then set the local parameters for the second core to reflect this request, undo the changes to the local parameters for the first core plus initiate changes to the domain parameters to transition the domain to the P1 state (the highest requested power/performance level of the two cores).

- l. In the paragraph that begins, "A hardware-independent dependency domain...", replace the following text:

A hardware-independent dependency domain relies on the software...

with:

A hardware-independent dependency domain (HIDD) relies on the software...

- m. Just after the paragraph that begins, "A hardware-independent dependency domain..." and just before the paragraph that begins, "The PAL procedure PAL_PROC_GET_FEATURES...", add the following text (including two section headers, to provide clarity):

11.6.1.2 Platform Power-Cap and P-states

Some processor implementations include mechanisms which allow the platform hardware and firmware to temporarily decrease the operating frequency of logical processors, to implement fast-response power capping. This is referred to as a Platform Power-Cap. In such implementations, the P-state requested by software is not changed by platform power-cap. Software is able to change its P-state request during platform power-caps, and when the platform power-cap is removed, the processor operating frequency returns to that of the P-state determined by software's most recent P-state settings.

The intention with platform power-caps is that they be of very short duration and very low duty cycle such that they do not have a significant effect on software-based methods for managing power through P-states. Platform power-caps do not affect the instantaneous operating P-state observed by software, but do affect the weighted-average performance index reported to software by PAL, so that software may take into account any small effects. (See the PAL_GET_PSTATE procedure for details.)

11.6.1.3 PAL Interfaces for P-states

- n. Replace the sentence:

The Itanium architecture provides three new PAL procedures to enable P-state functionality.

with:

The Itanium architecture provides three PAL procedures to enable P-state functionality.



- o. In the text portion that talks about “PAL_SET_PSTATE” and just before the paragraph that begins, “If the logical processor belongs to a software-coordinated dependency domain”, add the following text:

Implementation-specific event conditions may also cause a PAL_SET_PSTATE request to not be accepted. For example, thermal protection mechanisms to prevent over-temperature, if in effect, may cause PAL_SET_PSTATE to return a status of transition failure. These are expected to be rare, and to happen only in abnormal situations.

Note that platform power-caps do not cause PAL_SET_PSTATE requests to return status of transition failure. The newly requested P-state is registered with PAL, and the procedure returns a status of transition success.

- p. Replace the paragraph:

If the logical processor belongs to a software-coordinated dependency domain, the PAL_SET_PSTATE procedure will change the domain parameters, which will result in all logical processors in that domain to transition to the requested P-state. If the logical processor belongs to a hardware-coordinated dependency domain, the PAL_SET_PSTATE procedure will attempt to change the power/performance characteristics only for that logical processor, which will result in either partial or complete transition to the requested P-state. In case of partial transition (see Figure 11-22, “Computation of performance_index” on page 2:311 for an example, where the logical processor transitions from state P0 to state P3 in partial increments), the logical processor may attempt to perform changes at a later time to the local parameters and/or domain parameters to transition to the originally requested P-state. If the logical processor belongs to a hardware-independent dependency domain, the PAL_SET_PSTATE procedure will attempt to change the domain parameters, which will transition the logical processor in that domain to the requested P-state.

with:

SCDD: If the logical processor belongs to a software-coordinated dependency domain, the PAL_SET_PSTATE procedure will change the domain parameters, which will result in all logical processors in that domain to transition to the requested P-state.

HCDD: If the logical processor belongs to a hardware-coordinated dependency domain, the PAL_SET_PSTATE procedure will attempt to change the power/performance characteristics for that logical processor; since the power/performance characteristics for the domain depend on the P-state settings of the other logical processors in the domain, a PAL_SET_PSTATE call on one logical processor may result in either partial or complete transition to the requested P-state. In case of partial transition (see Figure 11-22, “Computation of performance_index” on page 2:311 for an example, where the logical processor transitions from state P0 to state P3 in partial increments), the logical processor may attempt to perform changes at a later time to the local parameters and/or domain parameters to transition to the originally requested P-state based on P-state transition requests on other logical processors. Software can also approximate the behavior of a SCDD by forcing P-state transitions. See the description of the PAL_SET_PSTATE procedure for more details.

HIDD: If the logical processor belongs to a hardware-independent dependency domain, the PAL_SET_PSTATE procedure will attempt to change the domain parameters, which will transition the logical processor in that domain to the requested P-state.

- q. In the paragraph that begins, “PAL_GET_PSTATE:...”, replace the last sentence:



Every invocation of the PAL_GET_PSTATE procedure resets the internal performance measurement logic, and initiates a new performance_index count, which is reported when the next PAL_GET_PSTATE procedure call is made.

with:

Software may choose, on each invocation of the PAL_GET_PSTATE procedure, whether to reset the internal performance measurement logic; resetting the measurement logic initiates a new performance_index count, which is reported when the next PAL_GET_PSTATE procedure call is made. A call to PAL_GET_PSTATE with a *type* operand of 1 resets the performance measurement logic.

- r. Replace the following paragraph:

If the logical processor belongs to a software-coordinated dependency domain or a hardware-independent dependency domain, the performance index returned corresponds to the target P-state requested by the most recent successful PAL_SET_PSTATE procedure call.

with:

SCDD: If the logical processor belongs to a software-coordinated dependency domain, the performance index returned (for either *type*=0 or 3) corresponds to the target P-state requested by the most recent successful PAL_SET_PSTATE procedure call. No weighted average (*type*=1 or 2) is computed by PAL; calling PAL_GET_PSTATE with *type*=1 or 2 on a SCDD logical processor is undefined.

- s. Replace the following sentence:

If the logical processor belongs to a hardware-coordinated dependency domain, the performance index returned will be a weighted-average sum of the perf_index values corresponding to the different P-states that the logical processor was operating in before the PAL_GET_PSTATE procedure was called.

with:


HCDD: If the logical processor belongs to a hardware-coordinated dependency domain, the performance index returned (*type*=1 or 2) will be a weighted-average sum of the performance_index values corresponding to the different P-states that the logical processor was operating in since performance measurement was last reset.

- t. Replace the following paragraph:

As seen above, for a hardware-coordinated dependency domain, the PAL_GET_PSTATE procedure allows the caller to get feedback on the dynamic performance of the processor over the last time period. The caller can use this information to get better system utilization over the next time period by changing the P-state in correlation with the current workload demand.

with:

As seen above, for a HCDD, the PAL_GET_PSTATE procedure allows the caller to get feedback on the dynamic performance of the processor over a software-controlled time period. The caller can use this information to get better system utilization over a subsequent time period by changing the P-state in correlation with the current workload demand. The caller can also use PAL_GET_PSTATE to see the most recent P-state set for this logical processor (*type*=0) and the instantaneous current P-state that the domain parameters are set to (*type*=3). Platform power-caps do not affect either of these return values.



HIDD: If the logical processor belongs to a hardware-independent dependency domain, a weighted-average performance index can be returned by PAL_GET_PSTATE (*type*=1 or 2). Since software could calculate on its own the performance index based on P-states set by software, the value of the weighted-average performance index is only in factoring in the effect of platform power-caps.

Note that P-state transitions typically do not happen instantaneously. An implementation-specific amount of time is required for a given transition to complete. The computation of the weighted-average performance_index may not take into account the fact that transitions of power/performance are gradual, but may be done as though they were instantaneous at the point when the transition starts. The expectation is that any errors in computing the performance_index due to non-instantaneous transitions to higher and lower P-states will tend to cancel out, and to the extent that they do not, will be insignificant.

- u. Just before the paragraph that begins “Some processors may support variable P-state performance...”, add this section heading:

11.6.1.4 Variable P-state Performance

2. Volume 2, Part I, Section 11.10.3, PAL_GET_PSTATE procedure:

- a. Replace the first three paragraphs of the description section:

This procedure returns the performance index of the processor over the time period between the previous and the current invocations of PAL_GET_PSTATE, and is relative to the highest available P-state, P0. A value of 100 represents the minimum processor performance in the P0 state. For processors that support variable P-state performance, it is possible for a processor to report a number greater than 100, representing that the processor is running at a performance level greater than the minimum P0 performance. The PAL procedure PAL_PROC_GET_FEATURES on page 2:429 indicates if the processor supports variable P-state performance.

For processors that belong to a software-coordinated dependency domain or a hardware-independent dependency domain, the *performance_index* value returned will correspond to the target P-state requested by the most recent PAL_SET_PSTATE procedure call in cases where variable P-state performance is not supported. When variable P-states performance is supported, the '*performance_index*' may be higher than the target P-state requested. Please see Section 11.6.1 for more information about variable P-state performance.

For processors that belong to a hardware-coordinated dependency domain, the *type* argument allows the caller to select the *performance_index* value that will be returned. See Table 11-73 below for details.


with:

This procedure returns a performance index of the processor, and is relative to the highest available P-state, P0. A value of 100 represents the minimum processor performance in the P0 state. For processors that support variable P-state performance, it is possible for a processor to report a number greater than 100, representing that the processor is running at a performance level greater than the minimum P0 performance. The PAL procedure PAL_PROC_GET_FEATURES on page 2:429 indicates if the processor supports variable P-state performance.

The *type* argument allows the caller to select the *performance_index* value that will be returned. See Table 11-73 below for details.



- b. Replace the following description of *type* 0 in Table 11-73:
- The *performance_index* returned will correspond to the target P-state requested by the most recent PAL_SET_PSTATE procedure call.
- with:
- The *performance_index* returned will correspond to the target P-state requested by software.
- * For SCDD (software-coordinated dependency domain) logical processors, this is the P-state requested by the most recent PAL_SET_PSTATE procedure call made by any logical processor in the domain.
 - * For HCDD (hardware-coordinated dependency domain) or HIDD (hardware-independent dependency domain) logical processors, this is simply the P-state requested by the most recent PAL_SET_PSTATE procedure call on this logical processor.
- The value returned is not affected by platform power-caps.
- c. Replace the description of *type* 3 in Table 11-73:
- The *performance_index* returned will correspond to the current instantaneous P-state of the logical processor, at the time of the procedure call.
- with:
- The *performance_index* returned will correspond to the current instantaneous P-state of the dependency domain containing the logical processor, at the time of the procedure call. The value returned is not affected by platform power-caps. When variable P-states performance is supported, the *performance_index* may be higher than the P-state requested. Please see Section 11.6.4 for more information about variable P-state performance.
- d. Replace the following paragraph:
- For processors that belong to a software-coordinated dependency domain or a hardware-independent dependency domain, the PAL_GET_PSTATE procedure should always be called with *type* argument value of 0 or 3. Note that the *performance_index* returned for *type*=0 and *type*=3 will have identical values for these coordination domains. This is because the most recent PAL_SET_PSTATE procedure call will always succeed in transitioning to the requested performance state for these coordination domains (see PAL_SET_PSTATE procedure description for additional details).
- with:
- For SCDD logical processors, or HIDD logical processors that do not support platform power-caps, note that the *performance_index* returned for *type*=0 and *type*=3 will have identical values. This is because the most recent PAL_SET_PSTATE procedure call that returned a *status* of 0 will always succeed in transitioning to the requested performance state for these coordination domains (see PAL_SET_PSTATE procedure description for additional details).
- For SCDD logical processors, the PAL_GET_PSTATE procedure should always be called with *type* argument value of 0 or 3. On such processors, calling PAL_GET_PSTATE with *type* argument value of 1 or 2 is undefined.
- For HIDD logical processors, the *type* argument values of 1 and 2 are supported, since such processors can also support platform power-caps, which affect the weighted-average performance index.
- e. Replace the following sentence:
- If there was a thermal-throttling event or any hardware-initiated event, which affected the processor power/performance for the current time period



and the accuracy of the *performance_index* value has been impacted by the event, then the procedure will return with *status*=1.

with:

If there was a thermal-throttling event or any hardware-initiated event other than platform power-caps which affected the processor power/performance for the current time period and the accuracy of the *performance_index* value has been impacted by the event, then the procedure will return with *status*=1.

f. Replace the following paragraph:

The procedure returns with a *performance_index* value of 100 when invoked for the first time. For subsequent invocations, the procedure will return the *performance_index* value corresponding to the processor performance in the time duration between the previous and current calls to PAL_GET_PSTATE.

with:

The procedure, when called with *type*=1 or *type*=2, returns with a fixed *performance_index* value of 100 until after the procedure has been called with *type*=1 to reset computation of the weighted-average *performance_index*. For subsequent invocations with *type*=1 or *type*=2, the procedure will return the *performance_index* value corresponding to the processor performance in the time duration between the previous call to PAL_GET_PSTATE with *type*=1 and the current call.

3. Volume 2, Part I, Section 11.10.3, PAL_SET_PSTATE procedure:

a. Replace the following text from the Description section:

PAL_SET_PSTATE is used to request the transition of the processor to the P-state specified by the *p_state* input parameter. The PAL_SET_PSTATE procedure does not wait for the transition to complete before returning back to the caller. The request may either be accepted (*status* = 0) or not accepted (*status* = 1), depending on hardware capabilities and implementation-specific event conditions. If the request is not accepted, then no transition is performed, and it is up to the caller to make another PAL_SET_PSTATE procedure call to transition to the desired P-state. When the request is accepted, it will attempt to initiate a transition to the requested performance state. For processors that belong to a software-coordinated dependency domain or a hardware-independent dependency domain, the procedure will always succeed in transitioning to the requested performance state. If the processor belongs to a hardware-coordinated dependency domain, the procedure will make a best-case attempt at fulfilling the transition request, based on the nature of the dependencies that exist between the logical processors in the domain. In such circumstances, the procedure may initiate no transition, partial transition or full transition to the requested P-state. Since there is the possibility that the procedure may initiate no processor transition, there are implementation-specific forward progress requirements.

The *force_pstate* argument may be used for a hardware-coordinated dependency domain when it is necessary to get a deterministic response for the P-state transition at the expense of compromising the power/performance of other logical processors in same domain. If the *force_pstate* argument is non-zero, and if the request is accepted, the procedure will initiate the P-state transition on the logical processor regardless of any dependencies that exist in the dependency domain at the time the procedure is called. The *force_pstate* argument is ignored for software-coordinated and hardware-independent dependency domain.



with:

PAL_SET_PSTATE is used to request the transition of the processor to the P-state specified by the *p_state* input parameter. The PAL_SET_PSTATE procedure does not wait for the transition to complete before returning back to the caller. The request may either be accepted (*status* = 0) or not accepted (*status* = 1), depending on hardware capabilities and implementation-specific event conditions. The presence of a platform power-cap does not prevent the request from being accepted. (See Volume 2, Section 11.6.1 for details.) If the request is not accepted, then no transition is performed, and it is up to the caller to make another PAL_SET_PSTATE procedure call to transition to the desired P-state. When the request is accepted, the processor will attempt to initiate a transition to the requested performance state. For SCDD or HIDD logical processors, the procedure will always succeed in transitioning to the requested performance state. For HCDD logical processors, the procedure will make a best-case attempt at fulfilling the transition request, based on the nature of the dependencies that exist between the logical processors in the domain. In such circumstances, the procedure may initiate no transition, partial transition or full transition to the requested P-state.

The *force_pstate* argument may be used for a HCDD when it is necessary to get a deterministic response for the P-state transition at the expense of compromising the power/performance of other logical processors in same domain. If the *force_pstate* argument is non-zero, and if the request is accepted, the procedure will initiate the P-state transition on the logical processor regardless of any dependencies that exist in the dependency domain at the time the procedure is called. Forcing the P-state does not change the P-states requested by other logical processors in the dependency domain, nor the value seen on other logical processors when they do a PAL_GET_PSTATE with *type*=0; rather, forcing the P-state effectively suspends hardware coordination. A subsequent call to PAL_SET_PSTATE on any logical processor in the dependency domain (with a *force_pstate* argument of zero) reinstates hardware coordination. The *force_pstate* argument is ignored on SCDD and HIDD logical processors.

4. Volume 2, Section 11.10.3, PAL_PSTATE_INFO procedure:

- a. In the bullet item describing *perf_index*, replace the sentence:

This field is enumerated on a scale of 0.100, with the value of 100 corresponding to the P0 state.

with:

This field is enumerated relative to the index of the highest-performing P-state, with the value of 100 corresponding to the minimum processor performance in the P0 state.

- b. In table 11-104, change the descriptions to read:

Hardware independent (HIDD)
Hardware coordinated (HCDD)
Software coordinated (SCDD)

19. Synchronization Requirements for Virtualization Opcode Optimization

This architectural change describes additional synchronization requirements for opcode optimization and makes the implicit SYNC_READ and SYNC_WRITE optional for PAL_VPS_SAVE and PAL_VPS_RESTORE respectively.

1. Volume 2, Part I, Section 11.7.4, "Virtualization Optimizations".

- a. Delete the last two sentences from the first paragraph:



Virtualization optimizations allow these instructions to execute, with `PSR.vm==1`, without causing intercepts to the VMM. Virtualization optimizations are divided into two classes:

b. Append the following to the end of first paragraph:

Virtualization optimizations reduce overall virtualization overhead, such as, for example, allowing these instructions to execute, with `PSR.vm==1`, without causing intercepts to the VMM. There are two types of virtualization optimizations - global and local. Local virtualization optimizations are further divided into virtualization accelerations and virtualization disables.

Global virtualization optimizations are specified during initialization of the virtual environment (that is, during `PAL_VP_INIT_ENV`). The specified optimizations are applicable to all the virtual processors running in the virtual environment. See Section 11.7.4.1, "Global Virtualization Optimizations" for details on the global virtualization optimizations supported in the architecture.

Local virtualization optimizations are specified during the creation of the virtual processor (i.e., during `PAL_VP_CREATE`). The optimization settings were specified in the VPD and hence local to each virtual processor. The VMM can specify different local optimization settings for different virtual processors. The two classes of local virtualization optimizations are:

2. Volume 2, Part I. Add a new Section 11.7.4.1, "Global Virtualization Optimizations":

11.7.4.1 Global Virtualization Optimizations

Table 11-19 summarizes the global virtualization optimizations supported in Itanium architecture.

Table 11-19. Global Virtualization Optimizations Summary

Optimization	<i>config_option</i> ^a	Description
Virtualization Opcode Optimization	opcode	Section 11.7.4.1.1
Virtualization Cause Optimization	cause	Section 11.7.4.1.2

Notes:

- a. *config_option* is a parameter for the `PAL_VP_INIT_ENV` procedure. See "PAL_VP_INIT_ENV – PAL Initialize Virtual Environment (268)" on page 2:462 for details.

For specific global virtualization optimizations, certain virtual processor control and architectural state is referenced directly by hardware/firmware, and hence must be maintained in the VPD, and synchronization is required when the VMM reads or writes this state in the VPD. Please refer to the corresponding section of each global virtualization optimizations for synchronization requirements.

11.7.4.1.1 Virtualization Opcode Optimization

Virtualization opcode optimization is enabled by the *opcode* bit in the *config_option* parameter of `PAL_VP_INIT_ENV`. Opcode information is provided to the VMM during PAL intercepts in the virtual environment. In some processor implementations, the opcode provided may not be guaranteed to be the opcode that triggered the intercept; virtual machine monitors can determine whether this is guaranteed from the *vp_env_info* return value of `PAL_VP_ENV_INFO`.

Table 11-20 and Table 11-14, "Virtual Processor Descriptor (VPD)" on page 2:314 shows the synchronization requirements and the VPD states that will be accessed for this optimization.



Table 11-20. Synchronization Requirements for Virtualization Opcode Optimization

VPD Resource	Synchronization Required
vpsr.ic	Write
vpsr.si	Write
vifa	Write
vitir	Write

11.7.4.1.2 Virtualization Cause Optimization

Virtualization cause optimization is enabled by the *cause* bit in the *config_option* parameter of *PAL_VP_INIT_ENV*. When enabled, the causes of virtualization intercepts will be provided to the VMM during PAL intercept handoffs within the virtual environment. When disabled, no cause information will be provided during PAL intercept handoffs.

This optimization requires no special synchronization.

3. Volume 2, Part I, Section 11.10.3, *PAL_VP_INIT_ENV* procedure.

a. In the fourth paragraph change the following:

Table 11-110 shows the layout of the *config_options* parameter. The *config_options* parameter configures the global configuration options for all the logical processors in the virtual environment. All logical processors in the virtual environment must specify the same configuration options in the *config_options* parameter, otherwise processor operation is undefined.

to:

Table 11-113 shows the layout of the *config_options* parameter. d The *config_options* parameter configures the global configuration options [and global virtualization optimizations] for all the logical processors in the virtual environment. All logical processors in the virtual environment must specify the same [value] in the *config_options* parameter [during *PAL_VP_INIT_ENV*], otherwise processor operation is undefined.

b. In the existing Table 11-110 "config_options - Global Configuration Options", add a new column on the left to classify bits into "Global Configuration Options" and "Global Virtualization Optimizations". (Note that Step 2 inserts two tables, renumbering Table 11-110 to Table 11-112.)




Table 11-112. *config_options* – Global Configuration Options

	Field	Bit	Description
Global Configuration Options	initialize	0	If 1, this procedure will initialize the PAL virtual environment buffer for this virtual environment. If 0, this procedure will not initialize the PAL virtual environment buffer. On a multiprocessor system, the VMM must wait until this procedure completes on the first logical processor before calling this procedure on additional logical processors; otherwise processor operation is undefined.
	fr_pmc	1	If 1, performance counters are frozen on all IVA-based interruptions when virtual processors are running. If 0, the performance counters will not be frozen on IVA-based interruptions when virtual processors are running.
	be	2	Big-endian – Indicates the endian setting of the VMM. If 1, the values in the VPD are stored in big-endian format and the PAL services calls are made with PSR.be bit equals to 1. If 0, the values in the VPD are stored in little-endian format and the PAL services calls are made with PSR.be bit equals to 0.
	Reserved	7:3	Reserved.
Global Virtualization Optimizations	opcode	8	If 1, opcode information will be provided to the VMM during PAL intercepts within the virtual environment. This opcode may or may not be guaranteed to be the opcode that triggered the intercept. See Table 11-111, “vp_env_info – Virtual Environment Information Parameter” on page 2:460 for details. If 0, most virtualization optimizations cannot be enabled through the virtualization acceleration control (<i>vac</i>) and virtualization disable control (<i>vdc</i>) fields in the VPD. For details on specific optimizations supported in <i>vac</i> and <i>vdc</i> , see Table 11-15, “Virtualization Acceleration Control (vac) Fields” on page 2:316 and Table 11-16, “Virtualization Disable Control (vdc) Fields” on page 2:317 . The value of this field also determines how virtualization events and General Exception faults are delivered to the VMM on certain instructions. See Section 11.7.2, “Interrupt Handling in a Virtual Environment” on page 2:317 for details.
	cause	9	If 1, the causes of virtualization intercepts will be provided to the VMM during PAL intercept handoffs within the virtual environment. No information will be provided if 0. If this field is 1, the <i>opcode</i> field also be 1, otherwise processor operation is undefined. See Section 11.7.3.1, “PAL Virtualization Intercept Handoff State” on page 2:320 for details of virtualization intercept handoffs.
	Reserved	63:10	Reserved.

4. Volume 2, Part I, Section 11.11, "PAL Virtualization Services".
 - a. PAL_VPS_RESTORE page. In the Arguments section, change the description of GR26-31 from "Scratch" to "Reserved".
 - b. PAL_VPS_SAVE page. In the Arguments section, change the description of GR26-31 from "Scratch" to "Reserved".
5. Volume 2, Part I, Section 11.11, "PAL Virtualization Services", PAL_VPS_SAVE page.
 - a. In the Arguments section, change the argument of GR26 from "Scratch" to "Skip implicit synchronization".
 - b. In the Description section, change the second paragraph from:

This service performs an implicit PAL_VPS_SYNC_READ; there is no need for the VMM to invoke PAL_VPS_SYNC_READ to synchronize the implementation-specific control resources before this service.

 to:



If GR26 is zero, this service performs an implicit PAL_VPS_SYNC_READ; there is no need for the VMM to invoke PAL_VPS_SYNC_READ to synchronize the implementation-specific control resources before this service. If GR26 is one (0x1), no implicit synchronization will be performed by this service.

6. Volume 2, Part I, Section 11.11, "PAL Virtualization Services", PAL_VPS_RESTORE page.

- a. In the Arguments section, change the argument of GR26 from "Scratch" to "Skip implicit synchronization".
- b. In the Description section, change the first sentence in second paragraph from:
This service performs an implicit PAL_VPS_SYNC_WRITE; there is no need for the VMM to invoke PAL_VPS_SYNC_WRITE unless the VPD values are modified before resuming the virtual processor.
to:

If GR26 is zero, this service performs an implicit PAL_VPS_SYNC_WRITE; there is no need for the VMM to invoke PAL_VPS_SYNC_WRITE unless the VPD values are modified before resuming the virtual processor. If GR26 is one (0x1), no implicit synchronization will be performed by this service.

20. New Priority Hint and Recommendations for Multi-Threading Hints

1. On the hint instruction page in Volume 3, update Table 2-31 as shown:

Table 2-31. Hint Immediates

<i>imm</i> ₂₁ or <i>imm</i> ₆₂	Mnemonic	Hint
0x0	@pause	Indicates to the processor that the currently executing stream is waiting, spinning, or performing low priority tasks. This hint can be used by the processor to allocate more resources or time to another executing stream on the same processor. For the case where the currently executing stream is spinning or otherwise waiting for a particular address in memory to change, an advanced load to that address should be done before executing a hint @pause; this hint can be used by the processor to resume normal allocation of resources or time to the currently executing stream at the point when some other stream stores to that address.
0x1	@priority	Indicates to the processor that the currently executing stream is doing a high priority task. This hint can be used by the processor to allocate more resources or time to this stream. Implementations will ensure that such increased allocation is only temporary, and that repeated use of this hint cannot impair longer-term fairness of allocation.
0x02-0x3f		These values are available for future architected extensions and will execute as a <i>nop</i> on all current processors. Use of these values may cause unexpected performance issues on future processors and should not be used.
<i>other</i>		Implementation specific. Performs an implementation-specific hint action. Consult processor model-specific documentation for details.

2. In Part II of Volume 1, add a new Section 4.5 after the current section 4.4 "Branch and Prefetch Hints":

4.5 Hints for Controlling Multi-threading

Some processors support multi-threading; that is, they support the simultaneous execution of multiple threads (multiple logical processors) through a common set of execution resources (data paths, functional units, TLBs, etc.). Functionally, each of these hardware threads fully implements the Itanium architecture; therefore, software need not be aware of multi-threading nor do anything special to support it. From



performance standpoint, there are a few circumstances where it may be beneficial for software to provide information about its future resource requirements, which can be done with the `hint` instruction. Such a hint could allow the processor to optimize resource allocation among the hardware threads.

Note that, although not all implementations support all types of `hint` instruction, those that do not support them execute the hint instruction as a `nop`, and hence there is little penalty for software to provide these hints.

4.5.1 Wait Loops

Say a thread is waiting for another software thread to complete a task and, during that time, doesn't expect to need significant processor resources but would like to receive its fair share of resources once the task is complete. In such a situation, the waiting thread can communicate this information to the processor as a hint. This encourages the processor to allocate more processor resources to other threads of execution while this thread is waiting.

Typically, the completion signal in question is a store, by some other software thread, to a particular memory location. For example, a software thread may be waiting to acquire a spinlock and may have little work to do until such time as it is able to acquire the lock. A store to the spinlock in question may be an indication that the lock is now available for this software thread to acquire.

This scenario can be hinted to the processor by executing an advanced load (`ld.a` or `ld.sa`) to the address that this software thread is waiting on, and then by executing a `hint @pause` instruction (in a subsequent instruction group). This encourages the processor to devote more resources to other threads, yet if an entry is invalidated from this thread's ALAT, normal processor resource allocation is resumed for this thread.

Resource allocation within the processor eventually reverts to a fair allocation, so there's no need for software to hint that it is no longer in a wait loop. Conversely, while software is in such a wait loop, it would be best to re-execute the `hint @pause` as part of that loop, to continue to assert the hint for as long as that thread is waiting.

Note that if there is some high likelihood that the ALAT may contain a large number of valid entries upon entering into a wait loop, there may be some advantage to removing these (e.g., with an `invala` instruction) prior to executing the advanced load to the address to be waited on. This may reduce the restoration of resource allocation to this thread in cases where ALAT entries get invalidated other than the one for the address being waited on, hence providing more processor resources to other threads.

4.5.2 Idle Loops

Another situation where a software thread expects not to need significant processor resources for the next little while is when the software thread is executing an OS-kernel idle loop. It can provide this information to the processor also by executing a `hint @pause` instruction. This encourages the processor to allocate more processor resources to other threads of execution for the next while.

Resource allocation within the processor eventually reverts to a fair allocation, so there's no need for software to hint that it is no longer in an idle loop. Conversely, while software is in such an idle loop, it would be best to re-execute the `hint @pause` as part of that loop, to continue to assert the hint for as long as that thread is idle.



Note that if there is some high likelihood that the ALAT may contain a large number of valid entries upon entering into an idle loop, there may be some advantage to removing these (e.g., with an `invalida` instruction) prior to entering the idle loop. This may reduce the restoration of resource allocation to this thread in cases where these ALAT entries get invalidated, hence providing more processor resources to other threads.

4.5.3 Critical Sections

The opposite case exists if software expects that, given extra resources for the next period of time, overall system performance and throughput would be optimized. For example, this software thread may be about to acquire a highly contested spinlock and enter a critical section of code, and expeditious progress through that critical section and the resultant speedy release of the spinlock may disproportionately benefit overall system performance and throughput.

This scenario can be hinted to the processor by executing a `hint @priority` instruction. This encourages the processor to devote more processor resources to this thread (at the expense of other threads) for some period of time.

Resource allocation within the processor eventually reverts to a fair allocation, so there's no need for software to hint that it is no longer in a critical section. Processors that support this hint also ensure that it cannot be abused to affect overall longer-term fairness of processor resource allocation.

4 Specification Clarifications

1. Clarification of `ptc.g` Release Semantics

1. Volume 3, Part I, Section 2.2, "Instruction Descriptions", `ptc.g` page: Change the following text in the Description section:

`ptc.g` has release semantics and is guaranteed to be made visible after all previous data memory accesses are made visible. The memory fence instruction forces all processors to complete the purge prior to any subsequent memory operations. Serialization is still required to observe the side-effects of a translation being removed.

to:

`ptc.g` has release semantics and is guaranteed to be made visible after all previous data memory accesses are made visible. Serialization is still required to observe the side-effects of a translation being removed. If it is desired that the `ptc.g` become visible before any subsequent data memory accesses are made visible, a memory fence instruction (`mf`) should be executed immediately following the `ptc.g`.
2. Volume 2, Part I, Section 4.4.7, "Sequentiality Attribute and Ordering". Change the following text in the fifth paragraph:

Global TLB purge instructions (`ptc.g` and `ptc.ga`) follow release semantics on the local processor as well as on remote processors, except with respect to global purge instructions being executed by that remote processor.

to:


Global TLB purge instructions (`ptc.g` and `ptc.ga`) follow release semantics on the local processor. They are also broadcast to all other processors in the TLB coherence domain; on each such remote processor, a point is chosen in its program-order execution and a local TLB purge operation is inserted at that point; this local TLB purge operation follows release semantics, except with respect to global purge instructions being executed by that remote processor.

2. Clarification of `PAL_MC_ERROR_INFO` Reporting of Uncacheable Transactions

1. In Volume 2, Part I, Section 11.3.2.1, "Processor State Parameter (GR 18)", Table 11-7. For the `cc` field, add the following statement at the end of the text in the description box:

This bit must not be set for non-cacheable transaction errors.
2. In Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications", `PAL_MC_ERROR_INFO` page:
 - a. In the paragraph beginning "**Cache_Check Return Format:**", add the following sentence immediately after the sentence ending: "...caches in the `level_index` input argument.":

The `cache_check` return format must be used to report errors in cacheable transactions. These errors may also be reported using the `bus_check` return format if the bus structures can detect these errors.
 - b. In the paragraph beginning "**Bus_check Return Format:**", add the following sentence immediately after the sentence ending: "...bus structure as specified in the `level_index` input argument.":



The `bus_check` return format must be used to report errors in uncacheable transactions. These errors must not be reported using the `cache_check` return format.

3. Clarification of behavior when `ptc.g` overlaps a translation register

1. Volume 2, Part II, Section 5.2.2.2.3, “`ptc.g`, `ptc.ga`”, fourth paragraph from the end, change the following:

The `ptc.g` instruction does not modify the page tables nor any other memory location. It affects both the local and all remote TC entries in the TLB coherence domain. It does not remove translations from either local or remote TR entries, and if a `ptc.g` overlaps a translation contained in a TR on either the local processor or on any remote processor in the coherence domain, the processor containing the overlapping translation will raise a Machine Check Abort.

to:

The `ptc.g` instruction does not modify the page tables nor any other memory location. It affects both the local and all remote TC entries in the TLB coherence domain. It does not remove translations from either local or remote TR entries. If a `ptc.g` overlaps a translation contained in a TR on the local processor, the local processor will raise a Machine Check Abort; if the `ptc.g` overlaps a translation contained in a TR on any remote processor in the coherence domain, no Machine Check Abort is raised.

4. INT3 Clarifications

1. Volume 2, Part I, Section 5.6, “Interruption Priorities”, Table 5-6. Change the following row:

80 IA-32 Breakpoint (INT 3) trap IA-32 Exception vector (Debug)

to:

80 IA-32 Breakpoint (INT 3) trap IA-32 Exception vector (Break)

2. Volume 2, Part I, Section 7.1, “Debugging”. In the “**Break Instruction fault**” bullet, change the following sentence:

Execution of the IA-32 INT 3 (break) instruction results in a `IA_32_Exception(Debug)` trap.

to:

Execution of the IA-32 INT 3 (break) instruction results in a `IA_32_Exception(Break)` trap.

5. Test feature instruction clarifications


1. Volume 1, Part I, Section 4.3.2, “Compare Instructions”.

- a. Change the first paragraph from:

Predicate registers are written by the following instructions: general register compare (`cmp`, `cmp4`), floating-point register compare (`fcmp`), test bit and test NaT (`tbit`, `tnat`), floating-point class (`fclass`), and floating-point reciprocal approximation and reciprocal square root approximation (`frcpa`, `fprcpa`, `frsqrrta`, `fprsqrrta`). Most of these compare instructions (all but `frcpa`, `fprcpa`, `frsqrrta` and `fprsqrrta`) set two predicate registers based on the outcome of the comparison. The setting of the two target registers is described below in Compare Types on page 1:53. Compare instructions are summarized in Table 4-8.

to:

Predicate registers are written by the following instructions: general register compare (`cmp`, `cmp4`), floating-point register compare (`fcmp`), test bit and test NaT (`tbit`, `tnat`), test feature (`tf`), floating-point class (`fclass`), and



floating-point reciprocal approximation and reciprocal square root approximation (*frcpa*, *fprcpa*, *frsqrrta*, *fprsqrrta*). Most of these compare instructions (all but *frcpa*, *fprcpa*, *frsqrrta* and *fprsqrrta*) set two predicate registers based on the outcome of the comparison. The setting of the two target registers is described below in Compare Types on page 1:53. Compare instructions are summarized in Table 4-8.

- b. Volume 1, Part I, Section 4.3.2, “Compare Instructions”, Table 4-8, add a new row, just under the row for *tnat*, with the following information:

Mnemonic	Operation
<i>tf</i>	Test feature

- c. Volume 1, Part I, Section 4.3.2, “Compare Instructions”, at the end of the following paragraph:

The test bit (*tbit*) instruction sets two predicate registers according to the state of a single bit in a general register (the position of the bit is specified by an immediate). The test NaT (*tnat*) instruction sets two predicate registers according to the state of the NaT bit corresponding to a general register.

add the following sentence:

The test feature (*tf*) instruction sets two predicate registers according to whether or not the selected feature is implemented in the processor.

- d. Volume 1, Part I, Section 4.3.3, “Compare Types”, Table 4-11, change the row for

tbit, *tnat*

to read:

tbit, *tnat*, *tf*

2. Volume 1, Part I, Section 3.4.2, “WAW Dependency Special Cases”, change the first sentence of the second paragraph from:

The set of compare-type instructions includes: *cmp*, *cmp4*, *tbit*, *tnat*, *fcmp*, *frsqrrta*, *frcpa*, and *fclass*.

to:

The set of compare-type instructions includes: *cmp*, *cmp4*, *tbit*, *tnat*, *tf*, *fcmp*, *frsqrrta*, *frcpa*, and *fclass*.

6. Clarification of Performance Counter Behavior Under Halt States

Volume 2, Part I, Section 7.2.3. “Performance Monitor Events”. Change the following text:

1. The number of retired instructions. These are defined as all instructions which execute without a fault, including nops and those which were predicated off.
2. The number of processor clock cycles the CPU is in either the NORMAL or LOW-POWER state (see Figure 11-19 on page 2:303).

to:

1. The number of retired instructions. These are defined as all instructions which execute without a fault, including nops and those which were predicated off.



Generic counters configured for this event count only when the processor is in the NORMAL or LOW-POWER state (see Figure 11-19 on page 2:303).

2. The number of processor clock cycles. Generic counters configured for this event count only when the processor is in the NORMAL or LOW-POWER state (see Figure 11-19 on page 2:303).

7. PMI Clarifications

1. Clarifications to Volume 2, Part I, Chapter 5, "Interrupts".

a. Volume 2, Part I, Section 5.1, "Interruption Definitions", change the following bullet:

- **Platform Management Interrupts (PMI)**

A platform management request to perform functions such as platform error handling, memory scrubbing, or power management has been received by a processor. The PALE_PMI entry point is entered to service the request. Program execution may be resumed at the point of interruption. PMIs are distinguished by unique vector numbers. Vectors 0 through 3 are available for platform firmware use and are present on every processor model. Vectors 4 and above are reserved for processor firmware use. The size of the vector space is model specific.

to:

- **Platform Management Interrupts (PMI)**

A platform management request to perform functions such as platform error handling, memory scrubbing, or power management has been received by a processor. The PALE_PMI entry point is entered to service the request. Program execution may be resumed at the point of interruption. PMIs are distinguished by unique vector numbers. Vectors 0 through 3 are available for platform firmware use and are present on every processor model. Vectors 4 through 15 are reserved for processor firmware use. See section 11.5, "Platform Management Interrupt (PMI)" for details.

b. Volume 2, Part I, Section 5.8.1, "Interrupt Vectors and Priorities", change the following paragraph:

PMIs have a separate vector space from external interrupts. PMI vectors 0-3 can be used by platform firmware. PMI vectors 4 and above are reserved for use by processor firmware. Assertion of the processor's PMI pin, when present, results in PMI vector number 0. PMI vector priorities are described in Chapter 11, "Processor Abstraction Layer."

to:

PMIs have a separate vector space from external interrupts. PMI vectors 0-3 can be used by platform firmware. PMI vectors 4 through 15 are reserved for use by processor firmware. Assertion of the processor's PMI pin, when present, results in PMI vector number 0. PMI vector priorities are described in "Platform Management Interrupt (PMI)".

c. Volume 2, Part I, Section 5.8.4.1, "Inter-processor Interrupt Messages", Table 5-17. At the end of the description for the PMI delivery mode:

PMI - pend a PMI interrupt for the specified vector to the processor listed in the destination. Allowed PMI vector values are 0-3. All other PMI vector values are reserved for use by processor firmware.

Add the following:

See Section 11.5, "Platform Management Interrupt (PMI)" for details.



2. Volume 2, Part I, Section 11.5.1, “PMI Overview”:
 - a. In the fifth paragraph (just above table 11-13), delete the last sentence:
“Vectors described as Intel reserved will be ignored by the processor.”
 - b. In Table 11-13, change the text of second column from the left from:
Intel Reserved PAL
to:
PAL Reserved
 - c. In Table 11-13, in the Description column, change both instances of
Intel Reserved
to:
PAL Reserved
3. Volume 2, Part I, Section 11.5.1 “PMI Overview”.
 - a. Change the first sentence of the first paragraph from:
PMI is an asynchronous highest-priority external interrupt that encapsulates...
to:
PMI is an asynchronous interrupt that encapsulates...
 - b. Change the second sentence of the third paragraph from:
PMI events are the highest priority external interrupts...
to:
PMI events are asynchronous interrupts higher priority than all external interrupts...

8. **PAL_MC_ERROR_INJECT Clarifications**

1. Volume 2, Part I, Section 11.10.3, “PAL Procedure Specifications”. In the PAL_MC_ERROR_INJECT procedure, change the 'err_data_buffer' argument description from:
64-bit physical address of a buffer providing additional parameters for the requested error. The address of this buffer must be 8-byte aligned.
to:
Unsigned 64-bit integer specifying the address of the buffer providing additional parameters for the requested error. The address of this buffer must be 8-byte aligned.
2. Volume 2, Part I, Section 11.10.3, “PAL Procedure Specifications”. In the PAL_MC_ERROR_INJECT specification, Table 11-95 “err_struct_info - Register File.”
 - a. For the “reg_num” field, change the “Bits” value from:
11:5
to:
12:5
 - b. For the “Reserved” field just under “reg_num” change the “Bits” value from:
31:12
to:
31:13



9. Min-state Save Area Clarifications

1. In Volume 2, Part I, Section 11.3.2.3, "Processor Min-state Save Area Layout". Add the following text, figure and table after Figure 11.14:

The NaT bits stored in the first entry of the min-state save area have the following layout.

Figure 11-15. Min-state Save Area NaT Bits




Table 11-9. Min-state Save Area Nat Bits Description

Bits	Description
0	Undefined (not used)
15:1	NaT bits for GR15 to GR1. Bit 1 represents GR1 and subsequent bits follow the ascending pattern
31:16	NaT bits for Bank 0 GR16 to GR31. Bit 16 represents Bank 0 GR16 and subsequent bits follow the ascending pattern.
47:32	NaT bits for Bank 1 GR16 to GR31. Bit 32 represents Bank 1 GR16 and subsequent bits follow the ascending pattern.
63:48	Undefined (not used).

2. Volume 2, Part I, Section 11.3.2, "PALE_CHECK Exit State". Change the first sentence from:
The state of the processor on exiting PALE_CHECK is:
to:
The state of the processor on exiting PALE_CHECK is listed below. For registers described as being saved to the min-state save area and available for use, the actual values in these registers are undefined unless specifically stated otherwise.
3. Volume 2, Part I, Section 11.4.2, "PALE_INIT Exit State". Change the first sentence from:
The state of the processor on exiting PALE_INIT is:
to:
The state of the processor on exiting PALE_INIT is listed below. For registers described as being saved to the min-state save area and available for use, the actual values in these registers are undefined unless specifically stated otherwise.

10. Semaphore Code Corrections

1. Volume 2, Part II, Section 2.4.1 "Spin Lock". In Figure 2-4 change the instruction:
`cmpxchg8.acq r1 = [lock], r2 ;; // attempt to grab lock`
to:
`cmpxchg8.acq r1 = [lock], r2, ar.ccv ;; // attempt to grab lock`
2. Volume 2, Part II, Section 2.4.3, "Dekker's Algorithm". In Figure 2-6, change the "cmp.eq" in the following code sequence from:
`ld8 r2 = [flag_you] ;;;; is other's flag 0?`
`cmp.eq p1, p0 = 0, r2`
`(p1) br.cond.spnt cs_skip ;;;; if not, resource in use`



to "cmp.ne" as shown:

```
ld8 r2 = [flag_you] ;/// is other's flag 0?  
cmp.eq p1, p0 = 0, r2  
(p1) br.cond.spnt cs_skip ;/// if not, resource in use
```



5 Documentation Changes

1. Revision 2.2 Documentation Changes

1. Volume 2, Part I, Section 8.3, "Interruption Vector Definition, Table 8-2. Change the second to last vector from:

External Interrupt vector

to:

Virtual External Interrupt vector

2. Volume 2, Part I, Section 8.3, "Interruption Vector Definition, Page 2:162, Table 8-4. Add the following vector to the table:

Vector Name	Offset	Page
Virtual External Interrupt vector	0x3400	2:177

3. Volume 2, Part I, Section 11.10.3, "PAL Procedure Specifications, PAL_PSTATE_INFO page, Figure 11-67. Change the name of the field {10:5} from "ddit" to "ddid".

§