

Arm® Memory System Resource Partitioning and Monitoring (MPAM) System Component Specification



Arm Memory System Resource Partitioning and Monitoring (MPAM) System Component Specification

Copyright © 2018-2024 Arm Limited or its affiliates. All rights reserved.

Release Information

The following releases of this document have been made.

Release history

Date	Issue	Confidentiality	Change
19 April 2024	A.a	Non-Confidential	First release of the MPAM system architecture supplement after the content was separated from <i>Memory System Resource Partitioning and Monitoring (MPAM), for A-profile Architecture, Arm Architecture Reference Manual Supplement, DDI0598</i> .

Proprietary Notice

This document is protected by copyright and other related rights and the use or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm Limited (“Arm”). **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether the subject matter of this document infringes any third party patents.

The content of this document is informational only. Any solutions presented herein are subject to changing conditions, information, scope, and data. This document was produced using reasonable efforts based on information available as of the date of issue of this document. The scope of information in this document may exceed that which Arm is required to provide, and such additional information is merely intended to further assist the recipient and does not represent Arm’s view of the scope of its obligations. You acknowledge and agree that you possess the necessary expertise in system security and functional safety and that you shall be solely responsible for compliance with all legal, regulatory, safety and security related requirements concerning your products, notwithstanding any information or support that may be provided by Arm herein. In addition, you are responsible for any applications which are used in conjunction with any Arm technology described in this document, and to minimize risks, adequate design and operating safeguards should be provided for by you.

This document may include technical inaccuracies or typographical errors. THE DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. Arm may make changes to the Document at any time and without notice. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, third party patents, copyrights, trade secrets, or other rights.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Reference by Arm to any third party’s products or services within this document is not an express or implied approval or endorsement of the use thereof.

This document consists solely of commercial items. You shall be responsible for ensuring that any permitted use, duplication, or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

This document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of this document shall prevail.

The validity, construction and performance of this License shall be governed by English Law.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. Please follow Arm's trademark usage guidelines at <https://www.arm.com/company/policies/trademarks>. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners.

Copyright © 2024 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

Arm document reference: PRE-20349

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by Arm and the party that Arm delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

The information in this manual is at EAC quality, which means that all features of the specification are described in the manual.

Web Address

<http://www.arm.com>

Contents

	Preface	
	About this book	viii
	Using this book	ix
	Conventions	xi
	Additional reading	xii
	Feedback	xiii
Chapter 1	Introduction	
1.1	Overview	1-15
1.2	Memory-system resource partitioning	1-16
1.3	Memory-system resource usage monitoring	1-17
1.4	Memory-system components	1-18
1.5	Versions of the MPAM System Component Architecture	1-19
1.6	Implementation flexibility	1-25
Chapter 2	MPAM and Arm Memory-System Architecture	
2.1	Overview	2-27
Chapter 3	ID Types, Properties, and Spaces	
3.1	Introduction	3-29
3.2	ID types and properties	3-30
3.3	PARTID spaces and properties	3-31
3.4	Maximum PARTID number	3-32
Chapter 4	Memory System Propagation of MPAM Information	
4.1	Introduction	4-34
4.2	Requester components	4-35
4.3	Terminating Completer components	4-36
4.4	Intermediate Completer-Requester components	4-37
4.5	Request buffering	4-38
4.6	Cache memory	4-39
4.7	MPAM for RME propagation of MPAM_SP with requests	4-40
Chapter 5	System Model	
5.1	Introduction	5-42
5.2	System-level field widths	5-44
5.3	Other Requesters with MPAM	5-45
5.4	Requesters without MPAM support	5-46

5.5	Model of a resource partitioning control	5-47
5.6	Interconnect behavior	5-48
5.7	Cache behavior	5-49
5.8	Memory-channel controller behavior	5-51
5.9	The MPAM for RME system	5-52
Chapter 6	MPAM in MSCs	
6.1	Introduction	6-60
6.2	Resource controls	6-61
6.3	Resource instance selection	6-62
6.4	Security in MSCs	6-68
6.5	Virtualization support in system MSCs	6-69
6.6	PE with integrated MSCs	6-70
6.7	System-wide PARTID and PMG widths	6-71
6.8	MPAM interrupts	6-72
6.9	MSC support of MPAM for RME	6-76
Chapter 7	Resource Partitioning Controls	
7.1	Introduction	7-79
7.2	MPAM partitionable resources	7-80
7.3	Standard partitioning control interfaces	7-81
7.4	Vendor or implementation-specific partitioning control interfaces	7-93
7.5	Measurements for controlling resource usage	7-94
7.6	PARTID narrowing	7-95
7.7	System reset of MPAM controls in MSCs	7-96
7.8	About the fixed-point fractional format	7-97
Chapter 8	Resource Monitors	
8.1	Introduction	8-100
8.2	MPAM resource monitors	8-101
8.3	Common features	8-104
8.4	Monitor configuration	8-108
8.5	Monitor behavior on overflow	8-109
Chapter 9	Memory-mapped Registers	
9.1	Overview of MMRs	9-113
9.2	Summary of memory-mapped registers	9-119
9.3	Memory-mapped ID register description	9-122
9.4	Memory-mapped partitioning configuration registers	9-164
9.5	Memory-mapped monitoring configuration registers	9-208
9.6	Memory-mapped control and status registers	9-269
Chapter 10	Errors in MSCs	
10.1	Introduction	10-288
10.2	Error conditions in accessing memory-mapped registers	10-289
10.3	Overwritten error status	10-293
10.4	Behavior of configuration reads and writes with errors	10-294
10.5	Optionality of error detection and reporting	10-299
Appendix A	Generic Resource Controls	
A.1	Introduction	A-301
A.2	Portion resource controls	A-302
A.3	Maximum-usage resource controls	A-303
A.4	Proportional resource allocation facilities	A-304
A.5	Combining resource controls	A-306
Appendix B	MSC Firmware Data	
B.1	Introduction	B-308

B.2	Partitioning-control parameters	B-309
B.3	Performance-monitoring parameters	B-310
B.4	Discovery of resource to RIS mapping	B-311
B.5	Discovery of wired interrupts	B-312

Preface

This preface introduces the MPAM System Component specification. It contains the following sections:

- *About this book.*
- *Using this book.*
- *Conventions.*
- *Additional reading.*
- *Feedback.*

About this book

This book is the *MPAM System Component Specification*.

It specifies:

- Memory-mapped registers and standard types of resource control interfaces for Memory-System Components, or MSCs.
- Memory-mapped registers and resource usage monitors for measuring resource usage in MSCs.

Together, these facilities permit software both to observe memory-system usage and to allocate resources to software by running that software in a memory-system partition.

This document defines the versions of the MPAM system component architecture. For more information, see [*Versions of the MPAM System Component Architecture*](#).

This document primarily describes hardware architecture. As such, it does not include information on either the software needed to control these facilities or the ways to implement effective controls of the memory system using the parameters defined by this architecture.

This document does not describe:

- The optional features to implement in a MSC.
- The resources and MSCs to be controlled by MPAM.

Intended audience

This document targets the following audience:

- Hardware and software developers interested in the MPAM system component architecture.

Using this book

This book is organized into the following chapters:

Chapter 1 *Introduction*

Read this chapter for an introduction to the MPAM extension.

Chapter 2 *MPAM and Arm Memory-System Architecture*

Read this chapter for a description of MPAM and Arm Memory-System Architecture.

Chapter 3 *ID Types, Properties, and Spaces*

Read this chapter for a description of ID Types, Properties, and Spaces.

Chapter 4 *Memory System Propagation of MPAM Information*

Read this chapter for a description of MSC Propagation of MPAM Information.

Chapter 5 *System Model*

Read this chapter for a description of the System model.

Chapter 6 *MPAM in MSCs*

Read this chapter for a description of MPAM in MSCs.

Chapter 7 *Resource Partitioning Controls*

Read this chapter for a description of Memory-System Partitioning.

Chapter 8 *Resource Monitors*

Read this chapter for a description of Performance Monitoring Groups.

Chapter 9 *Memory-mapped Registers*

Read this chapter for a description of Memory-Mapped Registers.

Chapter 10 *Errors in MSCs*

Read this chapter for a description of Errors in MSCs.

Appendix A *Generic Resource Controls*

Read this appendix for a description of Generic Resource Controls.

Appendix B *MSC Firmware Data*

Read this appendix for a description of MSC Firmware Data.

Glossary

Read this glossary for definitions of some of the terms that are used in this manual. The Arm Glossary does not contain terms that are industry standard unless the Arm meaning differs from the generally accepted meaning.

———— **Note** —————

Arm publishes a single glossary that relates to most Arm products, see the *Arm Glossary* through Arm Developer at <https://developer.arm.com/documentation/aeg0014/latest>. A definition in the glossary in this supplement might be more detailed than the corresponding definition in *Arm Glossary*.

How to read this book

For readers new to MPAM, read Chapters 1 to 5.

For readers interested in MPAM resource controls and memory-system component behaviors, read Chapters 6, 7, and Appendices A and B.

For readers interested in MPAM resource usage monitoring, read Chapters 8, 9, and 10.

For readers interested in pseudocode and pseudocode definitions, read the *Arm® Architecture Reference Manual for A-profile architecture* (ARM DDI 0487).

For readers interested in Realm Management Extension, RME, read the *Arm® Architecture Reference Manual for A-profile architecture* (ARM DDI 0487).

Conventions

The following sections describe conventions that this book can use:

- *Typographic conventions.*
- *Signals.*
- *Numbers.*
- *Pseudocode descriptions.*

Typographic conventions

The typographical conventions are:

italic Introduces special terminology, and denotes citations.

bold Denotes signal names, and is used for terms in descriptive lists, where appropriate.

monospace Used for assembler syntax descriptions, pseudocode, and source code examples.
Also used in the main text for instruction mnemonics and for references to other items appearing in assembler syntax descriptions, pseudocode, and source code examples.

SMALL CAPITALS

Used for a few terms that have specific technical meanings, and are included in the Glossary LINK.

Colored text Indicates a link. This can be:

- A URL, for example, <http://developer.arm.com>
- A cross-reference, that includes the page number of the referenced information if it is not on the current page, for example, [Signals](#).
- A link to a chapter or appendix, or to a glossary entry, or to the section of the document that defines the colored term.

Signals

In general this specification does not define processor signals, but it does include some signal examples and recommendations.

The signal conventions are:

Signal level The level of an asserted signal depends on whether the signal is active-HIGH or active-LOW. Asserted means:

- HIGH for active-HIGH signals.
- LOW for active-LOW signals.

Lower-case n At the start or end of a signal name denotes an active-LOW signal.

Numbers

Numbers are normally written in decimal. Binary numbers are preceded by `0b`, and hexadecimal numbers by `0x`. In both cases, the prefix and the associated value are written in a monospace font, for example `0xFFFF0000`.

Additional reading

This section lists relevant publications from Arm and third parties.

See Arm Developer, <https://developer.arm.com>, for access to Arm documentation.

Arm publications

This book contains information that is specific to this product. See the following documents for other relevant information:

- *Arm® Architecture Reference Manual for A-profile architecture* (ARM DDI 0487).
- *Arm® CoreSight Architecture Specification v2.0* (ARM IHI 0029).
- *ARM® Generic Interrupt Controller Architecture Specification, GIC architecture version 3.0 and version 4.0* (ARM IHI 0069).
- *Arm® System Memory Management Unit Architecture Specification, SMMU architecture version 3* (ARM IHI 0070).
- *The Realm Management Extension (RME), for SMMUv3 Arm® System Memory Management Unit Architecture Supplement* (ARM IHI 0094).

Feedback

Arm welcomes feedback on its documentation.

Feedback on this Manual

If you have any comments or queries about this Manual, create a ticket at <https://support.developer.arm.com>.

As part of the ticket, include:

- The title, *Arm® MPAM Memory System Component Specification*.
- The number, ARM IHI 0099A.a.
- The section name to which your comments refer.
- The page number(s) to which your comments refer.
- The rule identifier(s) to which your comments refer, if applicable.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

———— **Note** —————

Arm tests PDFs only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the appearance or behavior of any document when viewed with any other PDF reader.

Inclusive terminology commitment

Arm values inclusive communities. Arm recognizes that we and our industry have used terms that can be offensive. Arm strives to lead the industry and create change.

Previous issues of this document included language that can be offensive. We have replaced this language. To report offensive language in this document, email terms@arm.com.

Chapter 1

Introduction

This chapter contains the following sections:

- *Overview.*
- *Memory-system resource partitioning.*
- *Memory-system resource usage monitoring.*
- *Memory-system components.*
- *Versions of the MPAM System Component Architecture.*
- *Implementation flexibility.*

1.1 Overview

Computer systems running multiple applications or virtual machines (VMs) concurrently and on shared memory often have one or more of the following requirements:

- A requirement to control the performance effects of non-conforming software on the performance of other software.
- A requirement to bound the performance impact on software by other software.
- A requirement to minimize the performance impact of some software on other software.

These scenarios are common in enterprise networking and server systems. The Memory System Resource Partitioning and Monitoring (MPAM) architecture addresses these scenarios with two approaches that work together, under software control, to apportion the performance-giving resources of the memory system. The apportionment can be used to align the division of memory-system performance between software, to meet higher-level goals for dividing the performance of the system between software environments.

These approaches are:

- Memory-system resource partitioning.
- Memory-system resource usage monitoring.

The MPAM memory-system component architecture describes:

- Propagation of a Partition ID (PARTID) and Performance Monitoring Group (PMG) through the memory system.
- A framework for memory-system component (MSC) controls that partition one or more of the performance resources of the component. See [Memory-system components](#).
- An extension of the framework for MSCs to have performance monitoring that is sensitive to a combination of PARTID and PMG.
- Some implementation-independent, memory-mapped interfaces to memory-system component controls for performance resource controls most likely to be deployed in systems.
- Some implementation-independent memory-mapped interfaces to memory-system component resource monitoring that might be required to monitor the partitioning of memory-system resources.

Note

This specification documents the MPAM memory-system component (MSC) architecture. For more information about the MPAM PE architecture, see *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

There are different versions of this MPAM Extension. For more information, see: [Versions of the MPAM System Component Architecture](#).

1.2 Memory-system resource partitioning

The performance of programs running on a computer system is affected by the memory-system performance, which is in part controlled by several resources in the memory system. In a memory system shared by multiple virtual machines, operating systems, and applications, the resources available to one software environment can vary, depending on the other programs running that might consume more or less of an uncontrolled memory-system resource.

Memory-system resource partitioning provides controls on the limits and use of previously uncontrolled memory-system resources.

Shared, partitionable memory-system resources that can affect performance of a virtual machine, operating system, or application include:

- Shared caches, in which one application can displace the cached data of another application.
- Interconnect bandwidth, in which use by one application can interfere with use by another application due to contention for buffers, communication links, or other interconnect resources.
- Memory bandwidth, in which use by one application can interfere with the use by another application due to contention for DRAM bus bandwidth.

Memory-system performance resource partitioning is performed by MPAM resource controls located within the MSCs. Each memory-system component can implement zero or more MPAM resource controls within that component.

An MPAM resource control uses the PARTID that is set for one or more software environments. A PARTID for the current software environment labels each memory-system request. Each MPAM resource control has control settings for each PARTID. The PARTID in a request selects the control settings for that PARTID, which are then used to control the partitioning of the performance resources of that memory-system component.

1.3 Memory-system resource usage monitoring

Memory-system resource-usage monitoring measures memory-system resource usage. MSCs can have resource monitors. An MPAM monitor must be configured and enabled before it can be queried for resource-usage information. A monitor can be configured to be sensitive to a particular PARTID, or PARTID and PMG. Some monitors can be configured to certain subcategories of the resource, such as the memory bandwidth used by writes that use a PARTID and PMG.

A monitor can measure resource usage or capacity usage, depending on the resource. For example, a cache can have monitors for cache storage that measure the usage of the cache by a PARTID and PMG.

Monitors can serve several purposes. A memory-system resource monitor might be used to find software environments to partition, or a monitor's reads might be used to tune the memory-system partitioning controls. A PMG value can be used to subdivide the software environments within a PARTID for finer-grained monitoring results, or to make measurements over prospective partitions.

1.4 Memory-system components

A Memory-System Component (MSC) is a function, unit, or design block in a memory-system that can have partitionable resources. MSCs consist of all units that handle load or store requests issued by any MPAM Requester. These include cache memories, interconnects, Memory Management Units, memory channel controllers, queues, buffers, and rate adapters.

An MSC can be a part of another system component. For example, a PE can contain caches, which can contain MSCs.

———— **Note** ————

For more information about the MPAM PE architecture, see *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

1.5 Versions of the MPAM System Component Architecture

This document describes several versions of the MPAM system component architecture. The identification of architecture versions and the features present within a version are described in:

- *MPAM versions for MSCs.*
- *Relationships between MPAM versions.*
- *Interoperation of components with different MPAM versions.*

1.5.1 MPAM versions for MSCs

The architecture version of the MPAM Extension implemented in an MSC is described in the MPAMF_AIDR register fields, ArchMajorRev and ArchMinorRev. The MPAM Extension versions used in MSCs are a subset of the versions used in PEs as the MPAM MSC architecture does not cover the generation of MPAM information by MSCs that are not PEs. The architecture of the component specifies how that component generates MPAM information for memory-system requests that it originates.

MPAM Extension versions and the corresponding values of fields in MPAMF_AIDR of the MSC are shown in Table 1-1:

Table 1-1 MPAM version implemented by an MSC

MPAMF_AIDR		MPAM Extension version supported	MSC MPAM support
ArchMajor Rev	ArchMinor Rev		
0b0000	0b0000	None	The MSC does not implement MPAM.
0b0000	0b0001	n/a	Not a valid MPAM version for an MSC.
0b0001	0b0000	v1.0	The MSC implements MPAM v1.0 with features as described in the 32-bit MPAMF_IDR .
0b0001	0b0001	v1.1	The MSC implements MPAM v1.1 with features as described in the 64-bit MPAMF_IDR . MPAM v1.1 includes all of the MSC MPAM features of MPAM v1.0 plus additional MPAM features.

Most MPAM features in an MSC are optional. The particular MPAM features available in an MSC are described in the [MPAMF_IDR](#) register.

[MPAMF_IDR](#) is 32 bits in MPAM v1.0 and is 64 bits in MPAM v1.1.

[MPAMF_IDR](#) is permitted to have different MPAM features in different address spaces. If the MPAM feature RIS is implemented [MPAMF_IDR](#) is also permitted to have different features for different Resource Instances in an MSC.

MSCs can be used in MPAM v1.0 and v1.1, and in v0.1 under certain conditions. For more information on the conditions on use of MSCs in MPAM v0.1, see *MPAM versions in MSCs*.

If an MSC does not implement any of the MPAM v1.1 MSC features listed in *MPAM versions for MSCs*, then the MSC is of MPAM v1.0.

1.5.1.1 MSC of MPAM v1.1

The MPAM features that can be implemented in an MSC of MPAM v1.1 are:

Expansion of MPAMF_IDR

`MPAMF_IDR` is expanded to 64 bits to support bits that indicate the presence of features added from MPAM v1.1.

This feature is mandatory when the MSC implements MPAM v1.1.

This feature is implemented when `MPAMF_IDR.EXT` is set to 1.

For more information, see *MPAMF_IDR, MPAM Features Identification Register*.

Capturing of IMPLEMENTATION DEFINED resource partitioning controls or resource monitoring

This feature defines two fields that allow discovery of any IMPLEMENTATION DEFINED resource partitioning controls or IMPLEMENTATION DEFINED resource monitors that are implemented.

This feature is mandatory when the MSC implements MPAM v1.1 and `MPAMF_IDR.HAS_IMPL_IDR` is 1.

This feature is implemented when `MPAMF_IDR.EXT` is 1. Furthermore:

- When `MPAMF_IDR.NO_IMPL_PART` is 1, `MPAMF_IMPL_IDR` does not include the description of any implementation-specific resource partitioning controls.
- When `MPAMF_IDR.NO_IMPL_MSMON` is 1, `MPAMF_IMPL_IDR` does not include the description of any implementation-specific resource monitors.

For more information, see *MPAMF_IDR, MPAM Features Identification Register*.

Resource instance selection

Resource instance selection, or RIS, provides access to the control settings of multiple resources of the same type within one MSC.

This feature is optional when the MSC implements MPAM v1.1.

This feature is implemented when `MPAMF_IDR.EXT` and `MPAMF_IDR.HAS_RIS` are 1.

For more information, see

- *Resource instance selection.*
- *MPAMCFG_PART_SEL, MPAM Partition Configuration Selection Register.*
- *MSMON_CFG_MON_SEL, MPAM Monitor Instance Selection Register.*
- *Error conditions in accessing memory-mapped registers.*

Greater range for MBWU monitors

This feature supports 44-bit and 63-bit memory bandwidth usage counters.

This feature is optional when the MSC implements MPAM v1.1.

This feature is implemented when `MPAMF_MBWUMON_IDR.HAS_LONG` is 1.

For more information, see *Long MBWU counter and capture*.

Discovery of MPAMF_ESR and MPAMF_ECR

This feature supports the `MPAMF_IDR.HAS_ESR` field. This field indicates whether `MPAMF_ESR` and `MPAMF_ECR` are implemented.

This feature is mandatory when the MSC implements MPAM v1.1.

This feature is implemented when `MPAMF_IDR.EXT` is 1.

For more information, see *MPAMF_IDR, MPAM Features Identification Register*.

Expansion of MPAMF_ESR

This feature widens `MPAMF_ESR` to 64 bits to include space for a RIS field.

This feature is optional when the MSC implements MPAM v1.1. Implementation of this feature is mandatory if `MPAMF_IDR.{HAS_ESR, HAS_RIS}` are 1.

This feature is implemented when `MPAMF_IDR.{EXT, HAS_EXTD_ESR}` are 1.

For more information, see

- [MPAMF_ESR, MPAM Error Status Register](#)
- [Resource instance selection.](#)

1.5.2 MSC features by MPAM version

MPAM MSC features by MPAM version are shown in [Table 1-2](#):

Table 1-2 MSC features by MPAM version

MPAM feature	Subordinate feature	Subordinate feature 2	MPAM v1.0	MPAM v0.1/v1.1	MPAM for RME	ID field
Cache capacity partitioning			Optional	Optional	Optional	MPAMF_IDR.HAS_CCAP_PART
	Minimum cache capacity partitioning	-	Prohibited	Optional	Optional	MPAMF_CCAP_IDR.HAS_CMIN
	No maximum cache capacity partitioning	-	Prohibited	Optional	Optional	MPAMF_CCAP_IDR.NO_CMAX
	Cache maximum associativity partitioning	-	Prohibited	Optional	Optional	MPAMF_CCAP_IDR.HAS_CASSOC
	CMAX soft limit	-	Prohibited	Optional	Optional	MPAMF_CCAP_IDR.HAS_CMAX_SOFTLIM
	No maximum cache capacity partitioning	-	Prohibited	Optional	Optional	MPAMF_CCAP_IDR.NO_CMAX
Cache portion partitioning	-	-	Optional	Optional	Optional	MPAMF_IDR.HAS_CPOR_PART
PARTID disable			Prohibited	Optional	Optional	MPAMF_IDR.HAS_ENDIS
	No Future Use	-	Prohibited	Optional	Optional	MPAMF_IDR.HAS_NFU
Memory BW partitioning			Optional	Optional	Optional	MPAMF_IDR.HAS_MBW_PART
	Minimum BW partitioning	-	Optional	Optional	Optional	MPAMF_MBW_IDR.HAS_MIN
	Maximum BW partitioning	-	Optional	Optional	Optional	MPAMF_MBW_IDR.HAS_MAX
	Maximum BW partitioning limit behaviors	-	Optional	Optional	Optional	MPAMF_MBW_IDR.MAX_LIM
	BW portion partitioning	-	Optional	Optional	Optional	MPAMF_MBW_IDR.HAS_PBM
	Proportional BW partitioning	-	Optional	Optional	Optional	MPAMF_MBW_IDR.HAS_PROP
	BW window writable	-	Optional	Optional	Optional	MPAMF_MBW_IDR.WINDWR

Table 1-2 MSC features by MPAM version (continued)

MPAM feature	Subordinate feature	Subordinate feature 2	MPAM v1.0	MPAM v0.1/v1.1	MPAM for RME	ID field
Priority partitioning			Optional	Optional	Optional	MPAMF_IDR.HAS_PRI_PART
	Internal priority partitioning	-	Optional	Optional	Optional	MPAMF_PRI_IDR.HAS_INTPRI
	Downstream priority partitioning	-	Optional	Optional	Optional	MPAMF_PRI_IDR.HAS_DSPRI
Memory Sys resource monitoring			Optional	Optional	Optional	MPAMF_IDR.HAS_MSMON
	Cache storage usage monitoring	-	Optional	Optional	Optional	MPAMF_MSMON_IDR.MSMON_CSU
		CSU monitor capture	Optional	Optional	Optional	MPAMF_CSUMON_IDR.HAS_CAPTURE
		CSU monitor read-only	Optional	Optional	Optional	MPAMF_CSUMON_IDR.CSU_RO
		CSU monitor XCL	Prohibited	Optional	Optional	MPAMF_CSUMON_IDR.HAS_XCL
		CSU monitor overflow linkage	Prohibited	Optional	Optional	MPAMF_CSUMON_IDR.HAS_OFLOW_LNKG
		CSU monitor overflow status Reg	Prohibited	Optional	Optional	MPAMF_CSUMON_IDR.HAS_OFSR
		CSU monitor overflow capture	Prohibited	Optional	Optional	MPAMF_CSUMON_IDR.HAS_CEVNT_OFLOW
	Memory BW usage monitoring	-	Optional	Optional	Optional	MPAMF_MSMON_IDR.MSMON_MBWU
		MBWU monitor capture	Optional	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_CAPTURE
		MBWU monitor Long	Optional	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_LONG
		MBWU monitor R/W filtering	Optional	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_RWBW
		MBWU monitor scaling	Optional	Optional	Optional	MPAMF_MBWUMON_IDR.SCALE
MBWU monitor overflow linkage		Prohibited	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_OFLOW_LNKG	

Table 1-2 MSC features by MPAM version (continued)

MPAM feature	Subordinate feature	Subordinate feature 2	MPAM v1.0	MPAM v0.1/v1.1	MPAM for RME	ID field
		MBWU monitor overflow status Reg	Prohibited	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_OFSR
		MBWU monitor overflow capture	Prohibited	Optional	Optional	MPAMF_MBWUMON_IDR.HAS_CAPTURE
	Monitor overflow status register	-	Prohibited	Optional	Optional	MPAMF_MSMON_IDR.HAS_OFLOW_SR
	Monitor overflow MSI	-	Prohibited	Optional	Optional	MPAMF_MSMON_IDR.HAS_OFLW_MSI
	No hardwired overflow interrupt	-	Prohibited	Optional	Optional	MPAMF_MSMON_IDR.NO_OFLW_INTR
	Local monitor capture event generator	-	Optional	Optional	Optional	MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT
PARTID narrowing	-	-	Optional	Optional	Optional	MPAMF_IDR.HAS_PARTID_NRW
Implementation -defined ID Reg	-	-	Optional	Optional	Optional	MPAMF_IDR.HAS_IMPL_IDR
	Impl IDR no partitioning	-	Prohibited	Required	Required	MPAMF_IDR.NO_IMPL_PART
	Impl IDR no monitoring	-	Prohibited	Required	Required	MPAMF_IDR.NO_IMPL_MSMON
Extended ID register	-	-	Prohibited	Required	Required	MPAMF_IDR.EXT
Resource instance selector	-	-	Prohibited	Optional	Optional	MPAMF_IDR.HAS_RIS
Error status register			Prohibited	Optional	Optional	MPAMF_IDR.HAS_ESR
	Extended error status register	-	Prohibited	Optional	Optional	MPAMF_IDR.HAS_EXTD_ESR
Error MSI	-	-	Prohibited	Optional	Optional	MPAMF_IDR.HAS_ERR_MSI
Four PARTID spaces	-	-			Required	MPAMF_IDR.SP4

1.5.3 Relationships between MPAM versions

This section describes the relationships between MPAM versions.

1.5.3.1 MPAM v0.1

An MPAM v0.1 PE implements any permitted subset of the features of MPAM v1.1 and also implements MPAM3_EL3.FORCE_NS. The FORCE_NS field cannot be present in any other MPAM version.

1.5.3.2 MPAM v1.0

MPAM v1.0 is the base version of MPAM. Unless explicitly defined, all features from MPAM v1.0 are present in the other versions of MPAM.

In a PE that implements MPAM v1.0, the MPAM features available (either Required or Optional) are described in the *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

In an MSC that implements MPAM v1.0, the MPAM features available (either Required or Optional) are described in Table 1-2.

1.5.3.3 MPAM v1.1

MPAM v1.1 adds features beyond the base version of MPAM. Unless explicitly removed, all features from MPAM v1.1 are present in MPAM v0.1 and in MPAM for RME.

In a PE that implements MPAM v1.1, the MPAM features available (either Required or Optional) are described in the *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

In an MSC that implements MPAM v1.1, the MPAM features available (either Required or Optional) are described in Table 1-2.

1.5.3.4 MPAM for RME

The MPAM for RME architecture supports the Realm Management Extension (RME) in systems, PEs and MSCs.

MPAM for RME requires MPAM v1.1 or higher.

In a PE that implements both RME and MPAM, MPAM for RME is required.

In a PE, MPAM for RME requires the MPAM feature ALTSP.

In a PE that implements MPAM for RME, the MPAM features available (either Required or Optional) are described in the *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

In an MSC that implements MPAM for RME, the MPAM features available (either Required or Optional) are described in Table 1-2.

An MPAM for RME implementation requires support for 4 PARTID spaces, see *MPAM for RME propagation of MPAM_SP with requests*.

1.5.4 Interoperation of components with different MPAM versions

Hardware must not prevent PEs that implement different versions of the MPAM architecture to coexist within a system. However, PEs that implement different versions of the MPAM architecture might cause software issues.

There is no required relationship between the MPAM architecture version of a PE and the MPAM architecture version of an MSC accessed by the PE.

———— Note ————

For more information about the MPAM PE architecture, see *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487).

1.6 Implementation flexibility

Memory-system partitioning, monitoring capabilities, and certain implementation parameters must be discoverable by software, and they must be used by software to adapt to the system hardware. Discovery of MPAM memory-system component topology is expected to be by means of firmware data such as Device Tree or ACPI interface. MPAM controls and parameters of MSCs are discoverable in memory-mapped ID registers.

The width of memory-system partitioning and monitoring values communicated through the system can be sized to the needs of the system. The costs can thereby be adjusted to meet the market requirements.

This document defines standard interfaces to some resource partitioning and monitoring features of MSCs. It does so by defining ID registers that expose implementation parameters and options. It also defines configuration registers that allow standard programming of these features while giving substantial implementation flexibility. In addition, this document also defines a mechanism that permits IMPLEMENTATION DEFINED partitioning and monitoring features that might introduce partitioning or monitoring in new ways or of new resource types.

Chapter 2

MPAM and Arm Memory-System Architecture

This chapter contains the section:

- *Overview.*

2.1 Overview

This section is *informative*.

MPAM partitioning of memory-system performance resources must not affect the memory behavior specified in the *Arm® Architecture Reference Manual for A-profile architecture*. The Arm memory model, as specified in that manual, must be followed in all of its particulars, including requirements for observation, coherence, caching, order, atomicity, endianness, alignment, memory types, and any other requirements defined in the Arm memory model. Furthermore, these requirements must also be met:

- When the MPAM information in multiple requests to an MSC are the same or are different, and whether those multiple requests come from a single requestor or from multiple requestors.
- For all MPAM memory-system component resource controls and configurations.
- When MPAM information stored with data accessed from caches is the same as, or different from, MPAM information in requests that access that data.

Chapter 3

ID Types, Properties, and Spaces

This chapter contains the following sections:

- *Introduction.*
- *ID types and properties.*
- *PARTID spaces and properties.*
- *Maximum PARTID number.*

3.1 Introduction

This chapter is *normative*.

MPAM operation is based on the MPAM information that Requesters include with requests made to the memory-system.

This chapter defines the components of that MPAM information bundle, which consists of:

- Partition ID space (PARTID space).
- Partition number.
- Performance monitoring group.

Together the Partition ID space and Partition number uniquely identify an MPAM resource partition.

The MPAM information bundle is used by each MPAM-controlled resource that is accessed in the handling of a request. The Partition ID space and Partition number select resource control parameters particular to the resource.

3.2 ID types and properties

A partition is identified by its partition ID space and its partition number.

Partition ID spaces are related to the Security states and the physical address spaces but distinct from them as described in section *PARTID spaces and properties*.

A partition number references a particular partition within a partition ID space. A partition number in one partition ID space does not reference the same partition as the same partition number in a different partition ID space. For example, partition number 5 in one partition ID space is not the same as partition number 5 in a different partition ID space.

The numerical value of a partition number has no inherent meaning. The partition ID space and partition number in a request to the memory-system are used to select resource control parameters in memory-system components involved in transporting, handling, and completing the request.

Each controlled resource of each memory-system component has resource control parameters. The resource control settings for a particular partition are independent of the settings for other resources, other memory-system components and other partitions.

An MPAM resource partition has a single property, the performance monitoring group. The performance monitoring group is used to provide an additional filter for MPAM resource usage monitors to monitor a subset of software using a single partition.

In this document, PARTID is used for the partition number fields in registers. In the MPAM information bundle that accompanies memory-system requests, MPAM_SP or MPAM_NS is used for the partition ID space as it is encoded on the bus. PMG is used for the performance monitoring group fields in registers and on the bus.

The architectural maximum width of a PARTID field is 16 bits.

The architectural maximum width of a PMG field is 8 bits.

3.3 PARTID spaces and properties

MPAM has multiple PARTID spaces to permit separate management of the partition numbers and partition resource configurations by environments that cannot be managed as a single PARTID space due to separation or trust concerns.

MPAM uses physical PARTID spaces to communicate between Requesters and other memory-system components. Partitions in physical PARTID spaces are used to select the resource control settings in those memory-system components. Those control settings regulate the resource usage in that memory-system component. See [Chapter 4 *Memory System Propagation of MPAM Information*](#).

3.4 Maximum PARTID number

Each component implements a maximum PARTID number in each PARTID space that it supports. Component types are MSC or other Requester.

The range of valid PARTIDs is 0 to the maximum PARTID, inclusive. The maximum values of a PARTID implemented by different MSCs need not be the same.

Each MSC has an MPAM identification register with which to discover the maximum PARTID implemented in each physical PARTID space. The maximum Non-secure PARTID supported by an MSC is indicated in its MPAMF_IDR.PARTID_MAX. The maximum Secure PARTID supported by an MSC is indicated in its MPAMF_SIDR.PARTID_MAX.

Software must avoid using PARTIDs that exceed the smallest maximum of any MSCs accessed because the behavior of an MSC accessed with an out-of-range PARTID is CONstrained UNPREDICTABLE as described in *System-wide PARTID and PMG widths*.

Chapter 4

Memory System Propagation of MPAM Information

This chapter contains the following sections:

- *Introduction.*
- *Requester components.*
- *Terminating Completer components.*
- *Intermediate Completer-Requester components.*
- *Request buffering.*
- *Cache memory.*

4.1 Introduction

This section is *normative*.

The MPAM information bundle is propagated through the memory system components, or MSCs, that have MPAM resource controls or monitoring. The MPAM information bundle is described in *Introduction*.

MPAM information propagates in the direction of requests from Requesters towards terminating Completer components. This is the downstream direction. The upstream direction is from Completers towards Requesters.

The propagation behavior in the memory system depends on the function of the part of the memory system. Each MSC must implement at least one of the following behaviors:

- *Requester components.*
- *Terminating Completer components.*
- *Intermediate Completer-Requester components.*
- *Request buffering.*
- *Cache memory.*
- *MPAM for RME propagation of MPAM_SP with requests.*

If an MSC has no downstream components that use MPAM information, the MSC is not required to propagate MPAM information.

4.2 Requester components

Requesters must label all requests to downstream MSCs with MPAM information.

A Requester must have a device-appropriate means of setting the MPAM information in the request:

- The PE must use the scheme described in *PE Generation of MPAM Information, Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487)
- This architecture does not specify a mechanism for determining the MPAM information for requests from a non-PE Requester. Arm recommends that non-PE Requesters needing to use MPAM facilities specify a mechanism for determining the MPAM information. This consists of the PARTID space (MPAM_NS or MPAM_SP), the partition number (PARTID), and the performance monitoring group (PMG) for the memory system requests that it initiates.
- *Arm System Memory Management Unit Architecture Specification, SMMU architecture versions 3.0, 3.1 and 3.2* specifies MPAM information generation on memory system accesses translated by the SMMU and accesses originated by the SMMU to its tables in memory.
- *Arm Generic Interrupt Controller Architecture Specification, GIC architecture version 3.0 and version 4.0* specifies MPAM information generation on memory system accesses originated by the GIC to its tables in memory.

If a Requester does not support MPAM, the system must arrange to supply a value for MPAM information required for the interface. If no other mechanism is available, then these values must be driven to a default value, whether they are in the Non-secure physical PARTID space, the Secure physical PARTID space, the Root PARTID space or the Realm PARTID space.

See also *Requesters without MPAM support*.

4.3 Terminating Completer components

A terminating Completer receives requests from upstream Requesters but does not communicate the requests to a downstream Completer. Instead, the terminating Completer services the requests. A terminating Completer does not forward MPAM information from a request. A terminating MSC is the edge of MPAM in a system.

A DRAM controller is a terminating Completer, even though it communicates with DRAM devices to complete the request. The DRAM devices do not support MPAM communication, so MPAM information is not forwarded to them. This might also happen elsewhere in a system where there is no downstream Completer that has MPAM support.

4.4 Intermediate Completer-Requester components

Intermediate MSCs have both one or more Completer interfaces and one or more Requester interfaces.

An intermediate component can route a request from an upstream Requester to one of its downstream Requester ports. When routing a request from upstream to downstream, the intermediate component passes the MPAM information unaltered to the downstream Requester port.

An intermediate component might terminate some requests from upstream locally without propagating the request to a downstream Requester port if the request is serviced locally.

4.5 Request buffering

Requests can be buffered in any MSC. A request that is buffered must retain its MPAM information.

4.6 Cache memory

A cache line must store the MPAM information of the request that caused its allocation. See [Cache behavior](#) for requirements on cache memory behavior.

4.7 MPAM for RME propagation of MPAM_SP with requests

MPAM_SP is 2 bits in an MPAM for RME four PARTID space region. See *Four-space region*.

MPAM_SP must be propagated to all components within a four-space region.

MPAM_SP must be propagated to all bridges connecting a four-space region to a two-space region. See *Two-space region* and *Systems with both two PARTID space and four PARTID space components*.

MPAM_SP must be propagated from all bridges connecting two-space regions to a four-space region.

Chapter 5

System Model

This chapter contains the following sections:

- *Introduction.*
- *System-level field widths.*
- *Other Requesters with MPAM.*
- *Requesters without MPAM support.*
- *Model of a resource partitioning control.*
- *Interconnect behavior.*
- *Cache behavior.*
- *Memory-channel controller behavior.*
- *The MPAM for RME system.*

5.1 Introduction

This section describes a model of system behavior that can support the MPAM features. In particular, it describes the behavior of Requesters, interconnects, caches, and memory controllers.

In this system model, a request:

- Begins at a Requester, such as an I/O Requester, DMA controller, or graphics processor:
 - MPAM information, consisting of the PARTID space (MPAM_NS or MPAM_SP), the partition number (PARTID), and the performance monitoring group (PMG), is transported with every request.
- Traverses non-cache nodes that might be a transport component (such as an interconnect), a bus resizer, or an asynchronous bridge.
- Might reach an MSC that contains or is a cache:
 - Caches sometimes generate a response (cache hit) and sometimes pass the request on (cache miss).
 - Caches could also allocate entries based on the request.
 - Caches must store the MPAM PARTID, PMG, and MPAM_NS associated with an allocation:
 - Needed for cache-storage usage monitoring.
 - Used during eviction to another cache.
 - Cache eviction must attach MPAM fields to the eviction request. The source for MPAM information on an eviction might depend on whether the eviction is to memory or to another cache. See [Eviction and Optional cache behaviors](#).
- Might proceed from a cache to a transport component, and to other caches or a memory-channel controller.
- Might result in a memory controller or other terminating Completer device responding to a request it receives.

Figure 5-1 shows a simplified system model for the downstream flow, in the direction of requests from Requesters to Completers. In this figure, all objects implement an MSC except the PEs, I/O Requesters, and I/O Completers. PEs generate MPAM information from MPAM state in their System registers. I/O Requesters typically get their MPAM information when their requests pass through an SMMU.

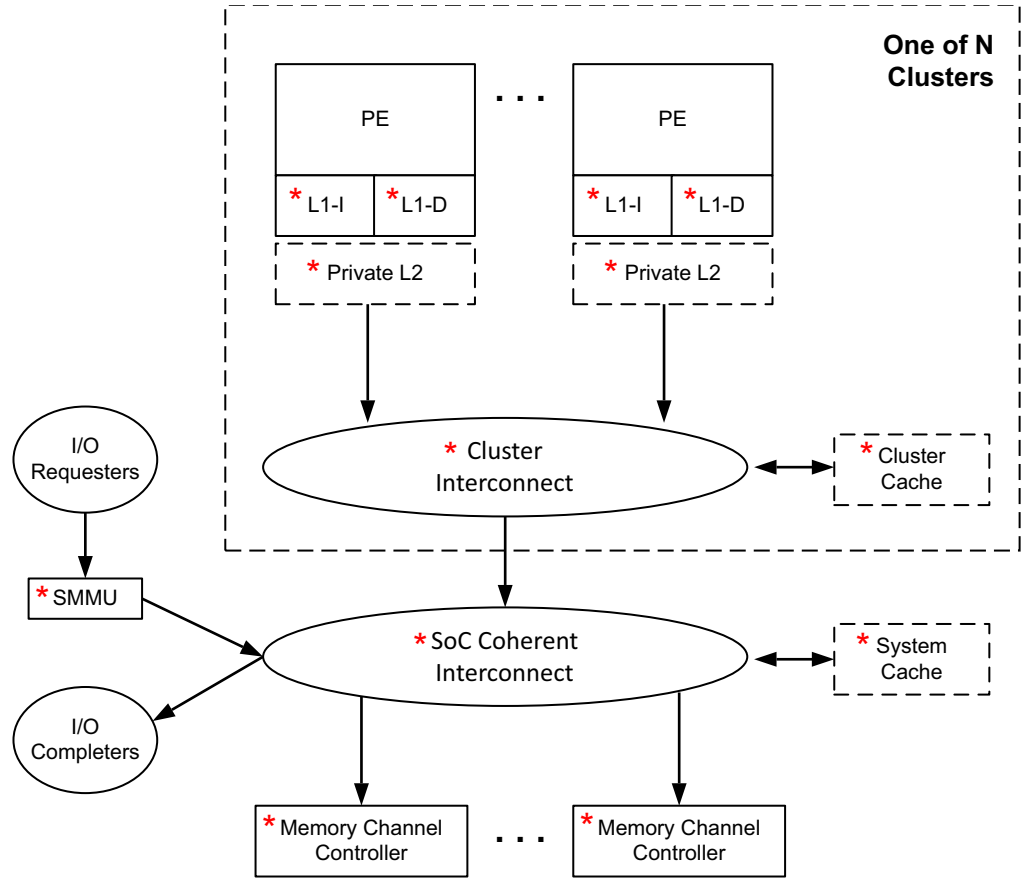
The interconnects in Figure 5-1 can represent bus, crossbar, packet, or other interconnect technologies.

An MSC responds to the MPAM information that arrives as part of a request. If the MSC implements partitioning controls, those controls find partitioning settings by the resource partition in the MPAM information of the request, and they use those settings to control the allocation of a controlled resource.

For caches, a cache line (which has an address) is always associated with the MPAM information of the request that allocated the line – or the MPAM information of the request that allocated the line into an inner cache that has now been evicted to the current cache. The inner cache PARTID must be preserved when the line is evicted to an outer cache.

An address can be accessed by multiple MPAM resource partitions.

A cache must store the MPAM information of the lines it contains, so that it can measure and control the cache lines used by a resource partition, and so that it can provide the MPAM information to downstream MSCs when the line is evicted.



* Memory-System Component (MSC) that might contain MPAM resource controls

Figure 5-1 MPAM system model (downstream flow)

5.2 System-level field widths

Arm recommends that a system be configured to support common values of maximum partition number (PARTID_MAX) and maximum PMG (PMG_MAX) in all Requesters and MSCs in the system.

Arm also recommends that when possible, the same PARTID spaces be supported throughout the system. See *The MPAM for RME system*.

5.3 Other Requesters with MPAM

Other Requesters that support MPAM, such as a DMA controller, must issue requests to the system that have the MPAM information. Non-PE Requesters can have schemes different from those implemented in PEs for associating MPAM information with requests. These other schemes are not documented in this specification.

5.4 Requesters without MPAM support

A Requester that does not implement support for MPAM must use a system-specific means to provide MPAM information to MSCs that support MPAM.

Some examples of Requester devices that might not implement support for MPAM include:

- Legacy DMA controller.
- Third-party peripheral IP.
- CoreSight DMA components, such as ETR.
- Older devices which cannot be economically upgraded to include MPAM support.

Some options for adding MPAM information to requests include:

- The MPAM information could be tied off to the default PARTID and PMG values and the PARTID space (MPAM_NS or MPAM_SP) set as appropriate for the device.
- The MPAM information could be provided by a System Memory Management Unit (SMMU) that supports adding MPAM information according to the stream and substream of the request.
- The MPAM information could be added by a bus bridge or other system component that handles the Requester's memory-system traffic.

Other implementations are permitted.

5.5 Model of a resource partitioning control

A general model of a resource partitioning controller within an MSC is shown in Figure 5-2. This model shows a resource partitioning model that measures resource usage by the partition and that controls resource usage by comparing the measured usage with the control settings for that partition.

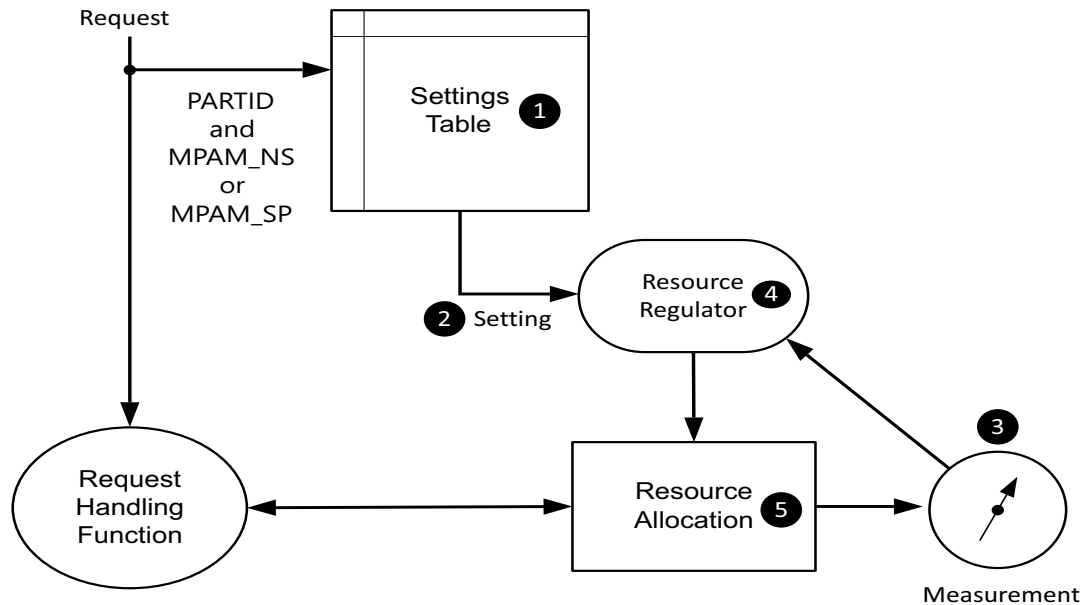


Figure 5-2 Model of MPAM resource partitioning controller

In Figure 5-2, a request arrives from an upstream Requester to an MSC that implements MPAM partitioning control. The request is handled as follows:

1. The PARTID space (MPAM_NS or MPAM_SP) and the partition number (PARTID) values of the incoming request are used to index into a Settings Table of partition-control settings. (There is one settings table per implemented resource control.)
2. The table entry for that resource partition specifies its partition-control setting, which is passed to a Resource Regulator.
3. Conformance of the resource with the setting might require measurement of how the resource is being used by the partition.
4. The measurement feeds back to the Resource Regulator, where it is compared with the Setting and used to make a decision about Resource Allocation.

In Figure 5-2, items 1, 2, 3, and 4 are added to the original memory system resource when MPAM resource control is implemented for the resource, although in some resources there might be sufficient measurement hardware already in place. Item 1, the Settings Table, is the heart MPAM resource control.

All of the above is separate from normal request-handling by the MSC.

Capacity-based partitioning requires the measurement of current usage of the resource by the partition as shown as item 3 above. The current resource usage measurement is compared to the resource control to determine whether the partition is using more or less than the setting.

Portion-based partitioning does not require the measurement in item 3 above as portions are predetermined and fixed.

5.6 Interconnect behavior

Interconnects connect Requesters to Completers, and they must transport MPAM information fields from Requester to Completer.

Interconnects can support the MPAM control features, such as priority partitioning. Support for MPAM is discoverable in ID registers and firmware data.

Some interconnect devices can include cache functionality, in which case the cache behavior in *Cache behavior* applies.

5.7 Cache behavior

A cache must associate the MPAM information of the request that allocated a cache line with any data stored in the cache line. This stored MPAM information is a property of the data.

The term “data” in this section is intended to indicate the content stored in the cache. It is not intended to indicate any restriction on the applicability of this section based on the purpose of the cache or of its content.

The MPAM information on a request to the cache from an upstream Requester is used for the following purposes:

- Source for the MPAM information associated with data when the data is allocated into the cache and is stored in association with the data while the data resides in the cache.
- Optionally updating the stored MPAM information of the cached data on a write hit (*Write hits may update the MPAM information of a cache line*).
- Providing MPAM information for downstream requests to fulfill the incoming request such as a read from downstream on a cache miss that fetches data into the cache.
- Optionally (*Eviction*), providing MPAM information for downstream requests generated by evict or clean operations when this cache is the last level of cache upstream of main memory.
- Selecting settings of partitioning controls implemented in the cache.
- Measuring or tracking the resource usage by each partition for a capacity control.
- Measuring or counting to track filtered resource usage for resource usage monitors, if implemented.
- Triggering and filtering events triggered by requests from upstream Requesters for MPAM resource monitors, if implemented.

The stored MPAM information is used by MPAM for the following purposes:

- Providing the MPAM information for downstream requests generated by evict or clean operations, when this cache is not the last level of cache.
- Optionally (*Eviction*) providing MPAM information for downstream requests generated by evict or clean operations, when this cache is the last level of cache.
- Triggering and filtering events triggered by internal and downstream requests for MPAM resource monitors, if implemented.
- Tracking resource usage by partitions, as needed by a partitioning control implementation.

5.7.1 Eviction

When a cache line is evicted to another cache, the evicting cache must produce the MPAM information that is associated with the cache line.

A system cache (last-level cache) might produce the MPAM information of the request that caused the eviction in its request to a memory-channel controller, or the cache might produce the stored MPAM information associated with the evicted line.

5.7.2 Cache partitioning

A cache may optionally implement cache-partitioning resource controls, such as a cache-portion partitioning control.

The cache-portion partitioning control (*Cache-portion partitioning*) was conceived for use on large, multi-way associative caches, but cache-portion partitioning can be implemented on caches that are not set-associative. For example, a single entry or group of entries may be a cache portion in a fully-associative cache.

The cache maximum-capacity partitioning control (*Cache maximum-capacity partitioning*) was conceived for use on caches that do not support cache-portion partitioning or that have insufficient portions to meet the needs of the planned use.

Both types of cache partitioning may be used together in a cache memory component. This may be useful, for example, when the cache has insufficient portions to give adequate control for a planned use.

5.7.3 Resource monitoring

A cache may implement cache-storage usage monitoring (*Cache-storage usage monitors*). For a monitored PARTID, the monitor gives the total cache storage used by the by partitions in the monitor's PARTID space and matching filter criteria programmed for the monitor. Filter criteria may include partition number (PARTID), performance monitoring group (PMG), and other criteria.

5.7.4 Optional cache behaviors

The following cache behaviors are permitted but not required.

5.7.4.1 Write hits may update the MPAM information of a cache line

If a write access, or a write-back from a cache level, associated with a request MPAM information, updates an entry that is present in the cache, it is CONstrained UNPREDICTABLE whether the cache entry stays associated with its stored MPAM information or is updated to be associated with the MPAM information of that write access.

It is possible that a change in the resource partition of the data (without moving the data) leaves the data in a portion of the cache that the new resource partition does not have permission to allocate. This can occur if the Cache Portion Bit Map (CPBM) bit for that portion is not set in the CPBM for the new PARTID. The optional behavior in this subsection does not change the location within the cache, even if the new partition for the data does not have a CPBM bit that allows allocation in this portion of the cache. Updating the location within the cache is a second optional behavior that is covered in the next subsection.

5.7.4.2 Write hits that update the resource partition of a cache line may move that line to a different portion

A write hit to cached data is permitted to change the portion of the cache capacity allocated to the data, if (i) the resource partition of the cache data is updated due to the write hit, and (ii) the portion of capacity where the data currently resides is not in the new resource partition's cache portion bitmap.

5.8 Memory-channel controller behavior

This section is *informative*.

A memory-channel controller may implement MPAM features. Some of the features that may be helpful in a memory-channel controller are:

- Memory-bandwidth minimum and maximum partitioning (*Memory-bandwidth minimum and maximum partitioning*).
- Memory-bandwidth portion partitioning (*Memory-bandwidth portion partitioning*).
- Priority partitioning (internal) (*Priority partitioning*).
- Memory-bandwidth usage monitors (*Memory-bandwidth usage monitors*).

5.9 The MPAM for RME system

5.9.1 Introduction

The MPAM for RME system supports RME_PEs and at least one PE that supports both RME and MPAM for RME.

RME_PEs support:

- Four Security states.
- Four physical address spaces.

A PE that supports RME and MPAM must also support MPAM for RME.

MPAM for RME requires support in the PE for:

- MPAM v1.1.
- Four MPAM PARTID spaces.
- MPAM *alternative space*, ALTSP feature.

There are three types of MPAM PARTID space regions that may be present in an MPAM for RME system. The regions are:

- Four-space regions.
- Two-space regions.
- Non-MPAM regions. See *Non-MPAM components*.

———— **Note** ————

The system must include a four-space region, but does not have to include two-space regions or non-MPAM regions.

It may be desirable to divide the system into regions that contain MSCs that support a like number of PARTID spaces and transport MPAM_SP or MPAM_NS at the width required to support those MSCs. A system may contain any number of regions of each type. Such regions can be labeled according to the number of PARTID spaces supported:

- Four-space regions transport MPAM_SP[1:0].
- Two-space regions transport MPAM_SP[0] or MPAM_NS.

Like other MPAM systems, MPAM for RME can also contain non-MPAM components and subsystems. See *Non-MPAM components*.

5.9.1.1 Four-space region

This type of region is distinguished by propagating MPAM information containing the 2-bit MPAM_SP:

- Contains one or more application PEs implementing FEAT_RME and FEAT_MPAM1p1.
- Contains caches associated with those PEs.
- Contains cache-coherent interconnect among those PEs that carry MPAM information containing the 2-bit MPAM_SP with requests.
- Contains only MSCs supporting the Non-secure PARTID space, the Realm PARTID space, the Root PARTID space and the Secure PARTID space.

All components in a four-space region must support and use four PARTID spaces.

If no Requester implements the Secure PARTID space because the Secure Security state is not implemented, the MPAM_SP encoding for the Secure PARTID space is unused and can be considered to be RESERVED.

5.9.1.2 Two-space region

This type of region contains a single two-space MPAM component or many two-space MPAM components connected as a subsystem through a two-space interconnect component. This component can connect to a four-space region using a bridging scheme.

Two-space MPAM components support two PARTID spaces. These are compatible with MPAM v1.0 and MPAM v1.1 but lack support for the Root and Realm PARTID spaces.

Two-space MPAM components can be used in an MPAM for RME system, but with some loss of functionality and with some complication to the MPAM software.

If a two-space region is within a system that has no Requesters supporting the Secure physical address space or the Secure PARTID space, the MPAM_NS encoding for the Secure PARTID space can be considered as RESERVED.

5.9.1.3 Systems with both two PARTID space and four PARTID space components

When two-space MPAM components are included in a four PARTID space system, all four-space MPAM components receive requests from any four PARTID space Requesters with all four states propagated to the four-space components.

If the propagation of the four PARTID spaces in the MPAM information labels is blocked by two-space components between any four-space Requester and any four-space Completers, the interface where the four PARTID spaces are reduced to two PARTID spaces is the boundary to a two-space region and must reduce the MPAM_SP to MPAM_NS using a bridge. The Completer is part of a two-space region and uses only two PARTID spaces even though it supports four.

Figure 5-3 shows an example of a system with a large four-space region with support for four PARTID spaces and a smaller two-space region. The boxes labeled 2 to 4 and 4 to 2 are bridges chosen from *Bridging between four-space and two-space regions*.

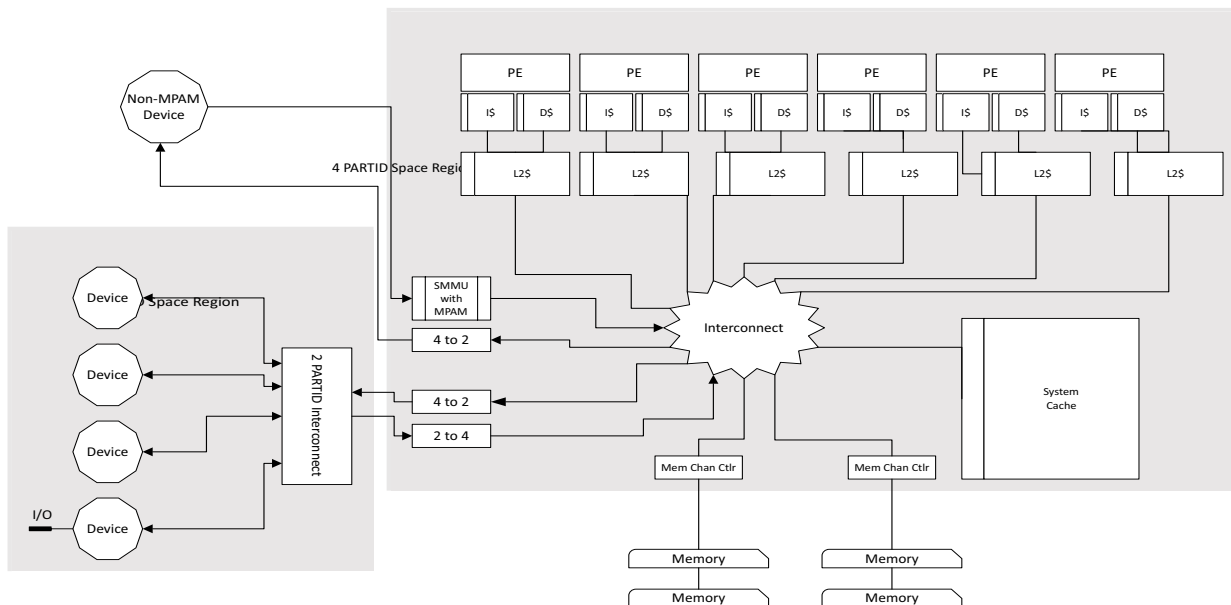


Figure 5-3 Example system with a large four PARTID space region and small two PARTID space regions

Figure 5-4 shows a system with a small four-space region and a large two-space region. In this case the bridges are not shown. Here the PEs can use the ALTSP feature to produce two PARTID space requests without the need for bridging logic. See *Fixed space mapping at a Completer*.

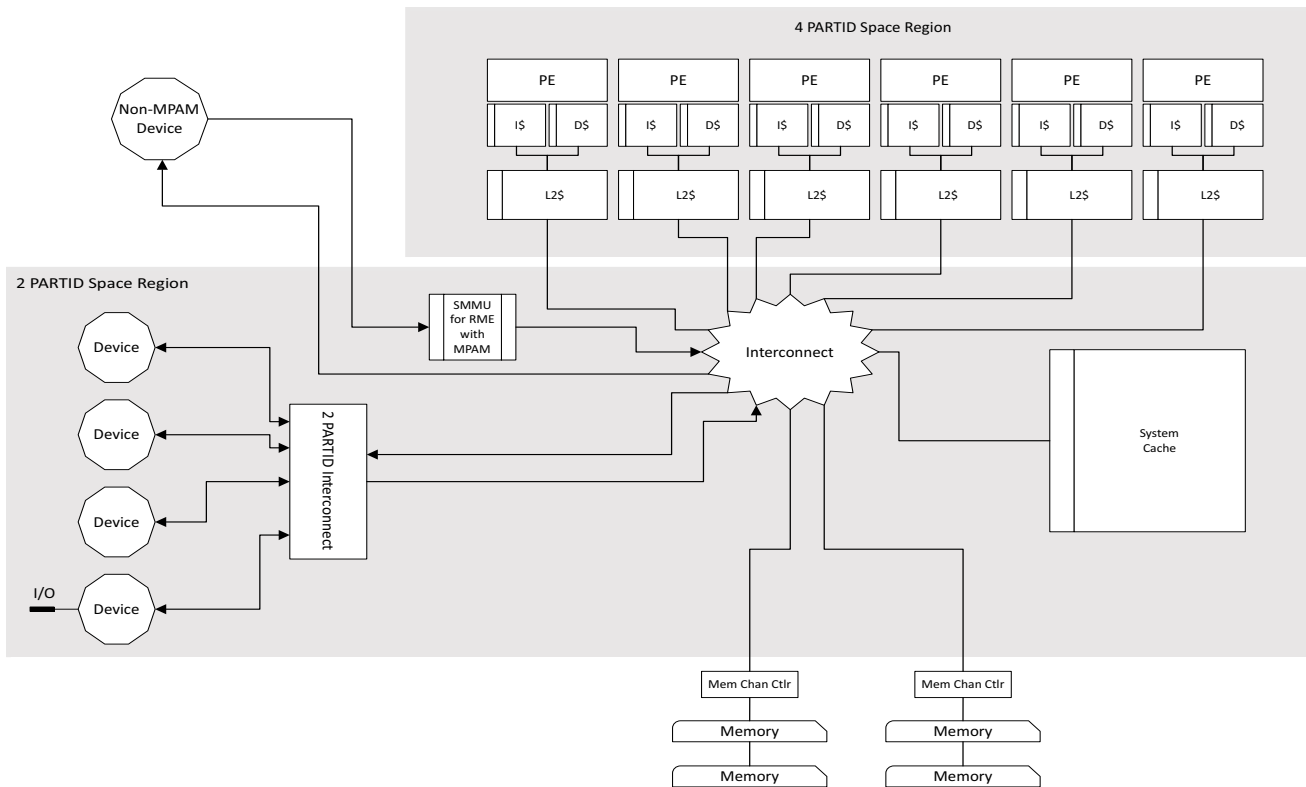


Figure 5-4 Example system with a small four PARTID space region and a large two PARTID space region

5.9.2 Requirements on bridges

The requirements on bridges are:

- The physical address space of a request must not be altered by bridging or other mechanisms.
- Bridging requests that use the Secure PARTID space must not be altered to use a different PARTID space.
- Bridging requests that use the Non-secure PARTID space must not be altered to use a different PARTID space.

5.9.3 Bridging between four-space and two-space regions

This section is *informative*.

Bridges are needed at the boundary between a four-space region and a two-space region. This section presents examples of bridging from two PARTID space Requesters to four PARTID space Completers and from four PARTID space Requesters to two PARTID space Completers. Bridging schemes other than the examples given in this section can also be implemented.

5.9.3.1 Two-Space Requesters

When a two-space MPAM Requester is upstream from a four-space MSC, the Requester's MPAM labels must have the MPAM_NS field expanded to the 2-bit MPAM_SP[1:0] while satisfying the requirements in [Requirements on bridges](#).

When bridging from a two-space region to a four-space region, Arm recommends a static mapping using the fixed MPAM_NS expansion.

5.9.3.1.1 Fixed MPAM_NS Expansion at a Requester

The fixed MPAM_NS expansion scheme transforms the MPAM_NS field to 2-bit MPAM_SP[1:0] field according to Table 5-1:

Table 5-1 Two-space Requester to four-space fixed expansion scheme

Two-space MPAM_NS Input	Four-space MPAM_SP[1:0] Output
0b0 (Secure PARTID space)	0b00 (Secure PARTID space)
0b1 (Non-secure PARTID space)	0b01 (Non-secure PARTID space)

The fixed expansion scheme preserves the PARTID space across the mapping.

5.9.3.2 Two-Space Completers

When a two-space MPAM Completer is downstream from a four-space Requester, the Requester's MPAM labels must have the MPAM_SP field reduced to form the 1-bit MPAM_NS while satisfying the requirements in *Requirements on bridges*. The reduction function may be static or dynamic.

———— **Note** —————

Arm makes no recommendation for which method to use for bridging between the four-space region of a system that has four PARTID spaces and a two-space region that supports two PARTID spaces. All known methods affect the system operation in ways that could cause difficulties for software.

5.9.3.2.1 Control over monitoring of Root and Realm PARTID space requests bridged to Secure or Non-secure PARTID space

A NO_MON flag is used in some of the examples to indicate that the transaction must not be monitored by MPAM monitors or other system performance monitors. This capability improves the security by limiting or preventing the system-level activities of a Realm from being collected in monitors accessible from the Non-secure physical address space or Secure physical address space.

The choice of not monitoring some transactions is not available on true two-space components. Support for the ability to mark requests with the NO_MON flag would likely require modifying the two-space component.

The examples that follow show a small number of recommended choices for including two-space MPAM Completer MSCs that do not have four-space MPAM support in an RME system. [Example 5-1](#) is the most desirable option, but requires extensive work in that it requires a redesign of the MSC. [Example 5-4](#) requires the least effort but is also the least desirable option:

Example 5-1 Alter the two-space MSC to support 4 PARTID spaces

This is the recommended option. However, it requires work to redesign the MSC. See *Four-space MSC* for how this is implemented.

Example 5-2 Alter the two-space MSC to support a programmable mapping of 4 PARTID spaces to 2 PARTID spaces

Alter the two-space MSC to support a programmable mapping of 4 PARTID spaces to 2 PARTID spaces with additional control over whether each of the Root and Realm PARTID spaces can be monitored. See *Programmable PARTID space mapping within a Completer*.

Example 5-3 Connect the two-space MSC through a programmable PARTID-space mapping component

Connect the two-space MSC through a programmable PARTID-space mapping component, or shim. See *Space mapping external to an MSC*.

This gives no control of whether the Root or Realm space can be monitored after being mapped into Secure or Non-secure.

Example 5-4 Connect the two-space MSC to be driven only from MPAM_SP[0]

Connect the two-space MSC so that the single-bit MPAM_NS input of the two-space MSC is driven only from MPAM_SP[0]. See *Fixed space mapping at a Completer*.

5.9.3.2.2 Programmable PARTID space mapping within a Completer

See [Example 5-2](#).

A programmable MPAM PARTID space mapping can be performed for a MSC with an PARTID space mapping built into the component. The PARTID space mapper accepts the request with 4 MPAM spaces, maps requests with MPAM_SP of Root or Realm to one of the Secure or Non-secure PARTID spaces and passes it on to the two-space MSC.

The programmable mapper can also produce a flag that indicates the two-space MSC should not perform MPAM monitoring of the request. See *Control over monitoring of Root and Realm PARTID space requests bridged to Secure or Non-secure PARTID space*.

The request mapper programming register is MAP4SPTO2SP. It has the fields shown in [Table 5-2](#):

Table 5-2 Request mapper programming register (MAP4SPTO2SP) fields

Field bits	Field name	Description
15	Rt_outPARTID_space	If a request has a Root PARTID, the output PARTID uses this bit for MPAM_NS.
14	Rt_NO_MON	If the request has a Root PARTID, output this bit as the NO_MON flag.
7	RI_outPARTID_space	If a request has a Realm PARTID, the output PARTID uses this bit for MPAM_NS.
6	RI_NO_MON	If the request has a Realm PARTID, output this bit as the NO_MON flag.

The MAP4SPTO2SP register must only be accessible in the Root physical address space.

5.9.3.2.3 Space mapping external to an MSC

See [Example 5-3](#).

A two-space Completer can be connected using a small component external to the MSC that implements a programmable four-space to two-space mapping similar to MAP4SPTO2SP. See Table 5-3:

Table 5-3 Space mapping external to the MSC MAP4SPTO2SP fields

Field bits	Field name	Description
15	Rt_outPARTID_space	If a request has a Root PARTID, the output PARTID uses this bit for MPAM_NS.
14	Rt_NO_MON	If the request has a Root PARTID, output this bit as the NO_MON flag.
7	RI_outPARTID_space	If a request has a Realm PARTID, the output PARTID uses this bit for MPAM_NS.
6	RI_NO_MON	If the request has a Realm PARTID, output this bit as the NO_MON flag.

The external mapping register must only be accessible in the Root physical address space.

If the two-space MSC does not have any way to accept the NO_MON flag at the request input, the NO_MON flag is not used. Two-space MSCs are not required to support a NO_MON input.

5.9.3.2.4 Fixed space mapping at a Completer

See Example 5-4.

The fixed MPAM_SP reduction scheme transforms MPAM_SP into a 1-bit MPAM_NS according to Table 5-4:

Table 5-4 Four-space to two-space static reduction scheme

Four-space MPAM_SP Input	Two-space MPAM_NS Output
0b00 (Secure PARTID space)	0b0 (Secure PARTID space)
0b01 (Non-secure PARTID space)	0b1 (Non-secure PARTID space)
0b10 (Root PARTID space)	0b0 (Secure PARTID space)
0b11 (Realm PARTID space)	0b1 (Non-secure PARTID space)

5.9.4 Non-MPAM components

Non-MPAM components do not have the ability to make requests with non-zero MPAM information or to use MPAM information when completing requests. They also do not propagate MPAM information to downstream MSCs.

5.9.4.1 Non-MPAM Requesters

Arm strongly recommends that a System MMU is used to add MPAM information to requests from non-MPAM Requesters, see *Arm® System Memory Management Unit Architecture Specification, SMMU architecture* (ARM IHI 0070) and *Arm® Realm Management Extension (RME), for SMMUv3 Arm® System Memory Management Unit Architecture Supplement* (ARM IHI 0094).

Requesters attached to an SMMU for RME are only associated with the Secure and Non-secure states, and therefore use two of the four PARTID spaces.

NoStreamID requesters attached to an SMMU for RME might issue transactions to Root or Realm physical address space. For these accesses it is permitted to use Secure and Non-secure PARTID spaces respectively.

5.9.4.2 Non-MPAM Completers

Completers that have no support for the MPAM information accompanying requests should be interfaced to the system by dropping MPAM information from the requests.

A non-MPAM Completer limits the topology of MPAM in the system because it does not propagate MPAM information to MPAM components downstream. See *Systems with both two PARTID space and four PARTID space components*.

Chapter 6

MPAM in MSCs

This chapter contains the following sections:

- *Introduction.*
- *Resource controls.*
- *Resource instance selection.*
- *Security in MSCs.*
- *Virtualization support in system MSCs.*
- *PE with integrated MSCs.*
- *System-wide PARTID and PMG widths.*
- *MPAM interrupts.*
- *MSC support of MPAM for RME.*

6.1 Introduction

This introduction to Memory-System Components, or MSCs, is *informative*. Other sections are normative unless marked as *informative*.

MSCs consist of all units that handle load or store requests issued by any MPAM Requester. These include cache memories, interconnects, Memory Management Units, memory channel controllers, queues, buffers, and rate adapters.

An MSC can be a part of another system component. For example, a PE might contain caches, which are MSCs. An MSC has resources that are used to process memory requests. The use of a resource can be controlled. A resource that can be controlled according to the PARTID of memory requests is partitioned. A resource might be monitored by a resource usage monitor.

6.1.1 MPAM versions in MSCs

MSCs can be used in MPAM v1.0, v1.1, and in v0.1 under certain conditions. If an MSC does not implement any of the MPAM v1.1 MSC features listed in *MPAM versions for PEs*, in *Arm® Architecture Reference Manual, for A-profile architecture* (ARM DDI 0487), then it is version 1.0.

———— **Note** —————

The MPAM version of an MSC is available in `MPAMF_AIDR`, see *MPAM versions for MSCs*.

If an MSC implements the extended `MPAMF_IDR`, or any of the MPAM v1.1 MSC features, it is either MPAM v1.1 or MPAM v0.1. An MSC must not use MPAM v0.1 unless all of these conditions are met:

- The MSC can initiate requests.
- Requests can be initiated in the Secure address space.
- Requests to the Secure address space can have `MPAM_NS` forced to 1.
- Software that configures the MSC to make requests in the Secure address space:
 - Cannot control the forcing of `MPAM_NS`.
 - Cannot easily see that `MPAM_NS` is being forced.

An MSC that supports the four physical address spaces of `FEAT_RME` must have `MPAMF_IDR.SP4` set to 1 and support an MPAM Feature page in each of the four address spaces. See *Four-space MSC*.

6.2 Resource controls

This section is *normative*.

An MSC optionally contains one or more MPAM resource controls. Although resource controls that control different performance resources have different control parameters, all resource controls are similar in the following aspects that form a common framework:

- Each resource control uses the MPAM PARTID and MPAM_NS signals from the incoming request to select control parameters from an array of Non-secure parameters (when MPAM_NS == 1) or Secure parameters (when MPAM_NS == 0).
- The selected parameters control the behavior of the MSC, either to partition the performance resources or to control the monitoring of performance resource usage.

For more information, see:

- [Model of a resource partitioning control](#) for a model of a resource partitioning control.
- [Chapter 7 Resource Partitioning Controls](#) for more detailed information on resource partitioning controls.
- [Resource instance selection](#) for how these controls are affected when resource instance selection is supported.

6.3 Resource instance selection

Resource instance selection, or RIS, allows support for MSCs with multiple resources. This includes multiple resources with the same resource type or partitioning control. This means that each MSC can only have independent resource controls and two or more resources of the same type when RIS is implemented. In MPAM v0.1 and from MPAM v1.1, this optional feature is implemented when `MPAMF_IDR.HAS_RIS` is 1.

This section provides more detail on:

- *RIS values.*
- *RIS controls in `MPAMCFG_PART_SEL`.*
- *RIS controls in `MSMON_CFG_MON_SEL`.*
- *Effects of `MPAMCFG_PART_SEL.RIS` on values read from other registers.*
- *Selecting a resource to monitor.*
- *Undefined RIS values.*
- *Reporting errors involving RIS.*

6.3.1 RIS values

Each resource that has MPAM resource partitioning controls or can be monitored by an MPAM resource usage monitor has a RIS value.

The RIS value in `MPAMCFG_PART_SEL.RIS` is used to select which resource to describe in ID register fields. `MPAMCFG_PART_SEL.RIS` is also used along with `MPAMCFG_PART_SEL.PART_SEL` to select the resource and PARTID when accessing `MPAMCFG_*` resource control registers.

MPAM resource monitors are usually associated with a resource instance, and the RIS value for that resource instance is also used in `MSMON_CFG_MON_SEL.RIS` to select the monitors associated with that resource.

RIS values are IMPLEMENTATION DEFINED. Any two resources in an MSC must have different RIS values. The RIS value is assigned to a resource in the MSC.

`MPAMF_IDR.RIS_MAX` gives the largest value of RIS that is defined for the MSC. A RIS value from 0 to `RIS_MAX` can be assigned to any partitioned or monitored resource. There is no requirement for every RIS value to be assigned to a partitioned or monitored resource.

As software for MPAMv1.0 would not set the value of the RIS field to any value other than 0, the only resource that can be identified and controlled by software that is not aware of this feature is resource instance 0.

6.3.2 RIS controls in `MPAMCFG_PART_SEL`

The value in `MPAMCFG_PART_SEL.RIS` selects the resource instance that is:

- Described by the MPAMF ID registers.
- Controlled by accessing the `MPAMCFG_*` registers.

6.3.3 Effects of `MPAMCFG_PART_SEL.RIS` on partitioning controls

To access control settings for a particular resource instance and PARTID, `MPAMCFG_PART_SEL.PART_SEL` is set to the PARTID and `MPAMCFG_PART_SEL.RIS` is set to the value associated with that resource instance. Accesses to additional `MPAMCFG_*` registers made without changing `MPAMCFG_PART_SEL` can be used to read and write additional control settings for that resource instance and partition.

If a control applies to all resource instances, this common control must be accessed with `MPAMCFG_PART_SEL.RIS` set to 0.

If there is only a single resource instance in an MSC, all controls must be associated with `MPAMCFG_PART_SEL.RIS` set to 0.

If an MPAMCFG register is accessed when `MPAMCFG_PART_SEL.RIS` is set to a resource instance that does not support the accessed control, then the behavior is `CONSTRAINED UNPREDICTABLE`, see *RIS in MPAMCFG_PART_SEL.RIS does not have partitioning control (errorcode == 9)*.

6.3.4 Effects of MPAMCFG_PART_SEL.RIS on values read from other registers

Fields within other registers reflect the capabilities of the resource instance that has been selected by `MPAMCFG_PART_SEL.RIS`, and so might have different values in different resource instances.

The effects of RIS on the MPAM identification registers are shown in Table 6-1

Table 6-1 MPAM ID register fields affected by a resource instance

Register	Field	Affected by MPAMCFG_PART_SEL.RIS
MPAMF_CCAP_IDR	CMAX_WD	This field is permitted to vary between resource instances.
	HAS_CMAX_SOFTLIM	This field is permitted to vary between resource instances.
	NO_CMAX	This field is permitted to vary between resource instances.
	HAS_CMIN	This field is permitted to vary between resource instances.
	HAS_CASSOC	This field is permitted to vary between resource instances.
	CASSOC_WD	This field is permitted to vary between resource instances.
MPAMF_CPOR_IDR	CPBM_WD	This field is permitted to vary between resource instances.
MPAMF_CSUMON_IDR	HAS_CAPTURE	This field is permitted to vary between resource instances.
	CSU_RO	This field is permitted to vary between resource instances.
	NUM_MON	This field is permitted to vary between resource instances.
	HAS_XCL	This field is permitted to vary between resource instances.
	HAS_CEVNT_OFLW	This field is permitted to vary between resource instances.
	HAS_CEVNT_CAPT	This field is permitted to vary between resource instances.
MPAMF_IDR	NO_IMPL_MSMON	MPAMF_IMPL_IDR describes no resource usage monitors.
	NO_IMPL_PART	MPAMF_IMPL_IDR describes no resource partitioning controls.
	HAS_MSMON	The resource usage monitors described in MPAMF_MSMON_IDR, otherwise this field is 0b0.
	HAS_IMPL_IDR	The IMPLEMENTATION DEFINED features described in MPAMF_IMPL_IDR, otherwise this field is 0b0.
	HAS_PRI_PART	The priority partitioning described in MPAMF_PRI_IDR, otherwise 0b0.
	HAS_MBW_PART	The memory bandwidth partitioning described in MPAMF_MBW_IDR, otherwise 0b0.
	HAS_CPOR_PART	The cache portion partitioning described in MPAMF_CPOR_IDR, otherwise 0b0.
	HAS_CCAP_PART	The cache capacity partitioning described in MPAMF_CCAP_IDR, otherwise 0b0.
MPAMF_IMPL_IDR	IMPLFEAT	The IMPLFEAT contents vary according to the resource instance selected, and cannot be specified by the architecture.
MPAMF_MSMON_IDR	MSMON_MBWU	The memory bandwidth usage monitors of the resource, otherwise this field is 0b0.
	MSMON_CSU	The cache storage usage monitors of the selected resource instance. Otherwise this field is 0b0.

Table 6-1 MPAM ID register fields affected by a resource instance (continued)

Register	Field	Affected by MPAMCFG_PART_SEL.RIS
MPAMF_PRI_IDR ^a	DSPRI_WD	The downstream priority width. Ignored if MPAMF_PRI_IDR.HAS_DSPRI is set to 0.
	DSPRI_0_IS_LOW	The downstream priority encoded with 0 being the low priority. Ignored if MPAMF_PRI_IDR.HAS_DSPRI is set to 0.
	HAS_DSPRI	The downstream priority control.
	INTPRI_WD	The internal priority width. Ignored if MPAMF_PRI_IDR.HAS_INTPRI is set to 0.
	INTPRI_0_IS_LOW	The internal priority encoded with 0 being low priority. Ignored if MPAMF_PRI_IDR.HAS_INTPRI is set to 0.
	HAS_INTPRI	The internal priority control.
MPAMF_MBW_IDR	BWPBM_WD	This field is permitted to vary between resource instances.
	HAS_PROP	This field is permitted to vary between resource instances.
	HAS_PBM	This field is permitted to vary between resource instances.
	HAS_MAX	This field is permitted to vary between resource instances.
	HAS_MIN	This field is permitted to vary between resource instances.
	HAS_MAX	This field is permitted to vary between resource instances.
	MAX_LIM	This field is permitted to vary between resource instances.
MPAMF_MBWUMON_IDR	HAS_CAPTURE	This field is permitted to vary between resource instances.
	HAS_RWBW	This field is permitted to vary between resource instances.
	HAS_LONG	This field is permitted to vary between resource instances.
	LWD	This field is permitted to vary between resource instances.
	SCALE	This field is permitted to vary between resource instances.
	NUM_MON	This field is permitted to vary between resource instances.
	HAS_OFLOW_LNKG	This field is permitted to vary between resource instances.
	HAS_OFSR	This field is permitted to vary between resource instances.
	HAS_CEVNT_OFLW	This field is permitted to vary between resource instances.
HAS_CEVNT_CAPT	This field is permitted to vary between resource instances.	

a. If the priority partitioning is local to the resource instance, then all fields might vary between resource instances. If the priority partitioning operates at the MSC level, then **MPAMF_PRI_IDR** should be non-zero only when **RIS** is 0.

The following registers are not affected by **RIS**:

- **MPAMF_AIDR.**
- **MPAMF_ECR.**
- **MPAMF_ESR.**
- **MPAMF_IIDR.**
- **MPAMF_PARTID_NRW_IDR.**
- **MPAMF_SIDR.**
- **MPAMCFG_PART_SEL.**
- **MPAMF_ERR_MSI_ADDR_H.**
- **MPAMF_ERR_MSI_ADDR_L.**

- MPAMF_ERR_MSI_ATTR.
- MPAMF_ERR_MSI_DATA.
- MPAMF_ERR_MSI_MPAM.

MPAMCFG resource control settings are selected by the MPAMCFG_PART_SEL register. The RIS field selects the resource instance and the PARTID_SEL field selects the PARTID of the resource control setting accessed. The PARTID space of the resource control is selected by the address space accessed.

MPAMCFG_PART_SEL.RIS accesses different resource instances. The resource instance selected can have control settings registers for accessing the controls for the selected resource.

The following resource configuration registers access resource control settings for different resources as selected by MPAMCFG_PART_SEL.RIS:

- MPAMCFG_CASSOC.
- MPAMCFG_CASSOC.
- MPAMCFG_CMIN.
- MPAMCFG_CPBM<n>.
- MPAMCFG_MBW_MAX.
- MPAMCFG_MBW_MIN.
- MPAMCFG_MBW_PBM<n>.
- MPAMCFG_MBW_PROP.
- MPAMCFG_PRI.

The following MPAM control settings are global to all resources in the MSC:

- MPAMCFG_DIS.
- MPAMCFG_EN.
- MPAMCFG_EN_FLAGS.
- MPAMCFG_INTPARTID.
- MPAMCFG_MBW_WINWD.
- MPAMCFG_PART_SEL.

6.3.5 RIS controls in MSMON_CFG_MON_SEL

The value in MSMON_CFG_MON_SEL.RIS selects the resource instance that is accessed by:

- The MSMON_CFG_* monitor configuration registers.
- The MSMON_* monitor and monitor capture registers.
- MSMON_CSU_OFSR and MSMON_MBWU_OFSR overflow status registers.

To access the configuration, value and capture registers associated with a monitor, the value of MSMON_CFG_MON_SEL.RIS should be set to match the RIS value associated with that monitor. Monitors not associated with any particular resource or associated with the MSC must be associated with MPAMCFG_PART_SEL.RIS == 0.

————— Note —————

Monitoring ID registers, MPAMF_MSMON_IDR, MPAMF_MBWUMON_IDR, and MPAMF_CSUMON_IDR, are not affected by MSMON_CFG_MON_SEL.RIS. These registers are affected by MPAMCFG_PART_SEL.RIS.

MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, MSMON_OFLOW_MSI_MPAM, and MSMON_OFLOW_SR are not affected by MSMON_CFG_MON_SEL.RIS.

6.3.6 Selecting a resource to monitor

To select the monitors for a particular resource instance, the value of MSMON_CFG_MON_SEL.RIS must be the same value as used in MPAMCFG_PART_SEL.RIS. Monitors that are not associated with an MPAM partitioned resource instance must be selected with a RIS value of 0.

To access a monitor for a particular resource, the MSMON_CFG_MON_SEL.RIS must be set to the resource instance. Then one or more MSMON_CFG_* registers for the particular monitor are accessed.

Any access to a MSMON_* register address will access the register associated with the resource instance value held in MSMON_CFG_MON_SEL.RIS. The exceptions to this are accesses to the MSMON_CFG_MON_SEL and MSMON_CAPT_EVNT registers, which are not affected by the value held in MSMON_CFG_MON_SEL.RIS.

6.3.7 Undefined RIS values

This section covers behaviors when the value of MPAMCFG_PART_SEL.RIS or MSMON_CFG_MON_SEL.RIS:

- Is greater than MPAMF_IDR.RIS_MAX.
- Does not correspond to an MPAM resource implemented in this MSC.
- Does correspond to an implemented MPAM resource, but the selected resource does not implement the control or monitor that has been accessed.

An MSC is permitted to:

- Implement fewer RIS bits than the architecture defines, though it must implement at least enough bits to represent MPAMF_IDR.RIS_MAX.
- Leave some RIS values that are within the range of 0 to MPAMF_IDR.RIS_MAX as undefined.
- Use only the implemented bits to decode RIS for selecting a resource instance.

Undefined resources that are within the range can still be identified. This is because the HAS_* fields within the ID registers all read as 0 when MPAMCFG_PART_SEL.RIS selects an undefined resource. All RIS values greater than MPAMF_IDR.RIS_MAX are undefined.

If software honors MPAMF_IDR.RIS_MAX and avoids accessing any *Memory-mapped registers* (MMR) that are not indicated with the corresponding HAS_* fields in the ID registers for that resource instance, it will not cause any RIS-related errors.

For more information on behavior caused by undefined RIS values, see:

- *Undefined RIS in MPAMCFG_PART_SEL.RIS (errorcode == 8).*
- *RIS in MPAMCFG_PART_SEL.RIS does not have partitioning control (errorcode == 9).*
- *Undefined RIS in MSMON_CFG_MON_SEL.RIS (errorcode == 10).*
- *RIS selected by MSMON_CFG_MON_SEL.RIS does not have monitor type (errorcode == 11).*

6.3.7.1 Reading an MPAMF ID register when MPAMCFG_PART_SEL is an undefined RIS value

Access to an MPAMF ID register when MPAMCFG_PART_SEL.RIS is an undefined value must produce an ID register value where all HAS_* fields read as 0. This action does not produce an error in MPAMF_ESR or signal an error interrupt.

6.3.8 Reporting errors involving RIS

Software can misconfigure the RIS fields in `MPAMCFG_PART_SEL` and `MSMON_CFG_MON_SEL` registers, possibly resulting in errors. See *Optionality of error detection and reporting*.

When an error is reported that involves a RIS value, the `MPAMF_ESR.RIS` field must be set to:

- For errors involving `MPAMCFG_*` register accesses, the `MPAMCFG_PART_SEL.RIS`.
- For errors involving `MSMON_*` register accesses, the `MSMON_CFG_MON_SEL.RIS` value.

For MPAM errors that do not capture the RIS field in `MPAMF_ESR.RIS` as shown in Table 10-1, `MPAMF_ESR.RIS` should be set to 0.

6.4 Security in MSCs

MPAM behavior in an MSC is affected in the following ways:

- Certain memory-mapped registers are only accessible from Secure address space (NS == 0).
- PARTIDs communicated to the MSC are augmented with a single MPAM_NS bit as 0, indicating that the MPAM PARTID in the request is to be interpreted in the Secure PARTID space. This is true even if the access from Secure state software was to the Non-secure (NS == 1) address space. MPAM_NS is always 0 if the PE is in the Secure state when the request is made, but the address of the request can be either a Secure or a Non-secure address. If the PE is in the Non-secure state, both the MPAM_NS bit and the address NS bit must be 1. See *PARTID spaces and properties*.
- When an MSC receives a transaction with MPAM_NS == 0, it accesses control settings for the Secure PARTID. If it receives a request with MPAM_NS == 1 it accesses the control settings for the Non-secure PARTID space.
- When programming the control settings for a Secure partition in an MSC, the settings must be stored by an access to the configuration registers in the Secure address space (NS == 0). See *Programming configuration of MPAM settings for Secure IDs*.
- When programming the control settings for a Non-secure partition in an MSC, the settings must be stored by an access to the configuration registers in the Non-secure address space (NS == 1).

6.4.1 Programming configuration of MPAM settings for Secure IDs

Configuration parameters for a Secure PARTID or Secure MPAM monitor can only be programmed from a Secure memory access (NS == 0):

- There are Secure and Non-secure versions of the MPAMCFG_PART_SEL and MSMON_CFG_MON_SEL. These two versions are accessed at the same address, differentiated by the value of the NS bit.
- Accessing an MPAMCFG_* register with a Secure (NS == 0) request accesses the configuration of a resource control of the Secure PARTID space that is selected by the PARTID in MPAMCFG_PART_SEL_S.
- Accessing an MPAMCFG_* register with a Non-secure (NS == 1) request accesses the configuration of a resource control of the Non-secure PARTID space that is selected by the PARTID in MPAMCFG_PART_SEL_NS.

6.4.2 Using Secure and Non-secure MPAM PARTIDs

When a request is processed by an MSC with MPAM resource controls, PARTID, PMG, and MPAM_NS control the partitioning control settings used and monitoring events triggered.

The PARTID and MPAM_NS of a request select the partitioning configuration from a table of PARTID configurations for each implemented resource control. The MPAM_NS bit in the request selects between the Non-secure configuration table and the Secure configuration table. The two tables do not need to have the same size. For example, the Secure configuration table might be much smaller. Tables are not required to be power-of-two sized.

A monitoring event is triggered if the PARTID, PMG, and MPAM_NS in a request match those configured in a performance monitor.

6.5 Virtualization support in system MSCs

MSCs do not see virtual PARTIDs. The PARTID generation in a Requester resolves any virtual PARTID into a physical PARTID that is communicated with the memory-system request. Therefore, MSCs only handle physical PARTIDs.

6.5.1 Hypervisor emulates guest accesses to partitioning and monitoring configurations

Accesses from a guest to the configuration registers of all MSCs, and to the System registers that configure the PE MSCs, may be emulated by the host hypervisor. This allows virtual PARTID mapping to be emulated and hypervisor policies governing resource partitioning to be applied.

Configuration and reconfiguration of control settings in MSCs are expected to be rare occurrences.

Arm recommends that the memory-mapped configuration registers of an MSC must be placed at a 64-KB-aligned address to permit an access trap on that page in the stage-2 page tables. The stage-2 access traps are taken to EL2 where the hypervisor can emulate the access. For more information on recommended configurations of memory-mapped registers of an MSC, see [MPAM feature page](#).

6.6 PE with integrated MSCs

A PE might have integrated MSC behaviors. These are discovered and configured in the same way as other MSCs.
See: [Chapter 9 Memory-mapped Registers](#) .

6.7 System-wide PARTID and PMG widths

This section is *informative*.

The behavior of an MSC is **CONSTRAINED UNPREDICTABLE** if it receives an MPAM PARTID or PMG outside the range it supports. For more information, see *Behavior of configuration reads and writes with errors*.

For predictable behavior, the PARTID on a request by a Requester must be in the range of 0 to:

- If the request is `MPAM_NS == 1`, the smallest maximum Non-secure PARTID supported by any MSC that might be accessed by that request.
- If the request is `MPAM_NS == 0`, the smallest maximum Secure PARTID supported by any MSC that might be accessed by that request.

And, the PMG on a request by a Requester must be in the range of 0 to:

- If the request is `MPAM_NS == 1`, the smallest maximum Non-secure PMG supported by any MSC that might be accessed by that request.
- If the request is `MPAM_NS == 0`, the smallest maximum Secure PMG supported by any MSC that might be accessed by that request.

The smallest maximum values for PARTID and PMG in Non-secure and Secure spaces can be computed from firmware during discovery. PARTID and PMG widths are reported through ID registers in PEs and MSCs. See sections *Appendix B MSC Firmware Data*, and *Determining presence and location of MMRs*.

6.8 MPAM interrupts

This section is *normative*.

There are two types of interrupts that an MPAM MSC can generate:

- MPAM Error Interrupt.
- MPAM Overflow Interrupt.

6.8.1 MPAM Error Interrupt

MPAM errors in MSCs are described in *Error conditions in accessing memory-mapped registers*.

MPAM errors that are detected in an MSC are recorded in `MPAMF_ESR` and signaled to software via an MPAM error interrupt if enabled by `MPAMF_ECR.INTEN == 1`.

If an MSC cannot encounter any of the error conditions listed in *Error conditions in accessing memory-mapped registers*, both the `MPAMF_ESR` and `MPAMF_ECR` must be RAZ/WI. An error cannot be encountered if the MSC:

- Does not support any feature of MPAM that can raise that error.
- Is designed so that the error cannot occur.
- Is permitted to have no detection for that error and does not implement detection for the error, see *Required error condition detection*.

If an MSC supports both Secure and Non-secure address spaces, `MPAMF_ESR` and `MPAMF_ECR` will each have a Secure instance and a Non-secure instance. The Secure registers control and generate Secure MPAM error interrupts, while the Non-secure registers control and generate Non-secure MPAM error interrupts.

The MPAM error interrupt can be implemented in an MSC as a level-sensitive interrupt or as an edge-triggered interrupt. The interrupt behavior depends on whether level-sensitive or edge-triggered interrupts are used.

- Arm recommends that the MPAM error interrupt be implemented as a level-sensitive interrupt.
- The mechanism by which an interrupt request from an MSC resource monitor generates an FIQ or IRQ exception is IMPLEMENTATION DEFINED.
- Arm recommends that an MSC implements two MPAM error interrupt signals, one for the Secure MPAM error interrupt and another for the Non-secure MPAM error interrupt.
- Arm recommends that MPAM error interrupt requests:
 - Translate into an `MPAM_ERR_IRQ` signal, so that they are observable to external devices.
 - If the MSC is integrated into a PE, connect to inputs on an IMPLEMENTATION DEFINED generic interrupt controller as a Private Peripheral Interrupt (PPI) or a Locality-specific Peripheral Interrupt (LPI) for that PE. See the *Arm Generic Interrupt Controller Architecture Specification, GIC architecture version 3.0 and version 4.0* for information about PPIs, LPIs, and SPIs.
 - If the MSC is not integrated into a PE, connect to inputs on an IMPLEMENTATION DEFINED generic interrupt controller as a System Peripheral Interrupt (SPI) or Locality-Specific Peripheral Interrupt (LPI).

6.8.1.1 Level-sensitive interrupts

When using level-sensitive interrupts, the interrupt is active when `MPAMF_ESR.ERRCODE` is non-zero.

Software can make a level-sensitive interrupt active by writing non-zero to `MPAMF_ESR.ERRCODE`.

An interrupt service routine is expected to write `0b0000` into `MPAMF_ESR.ERRCODE` to clear the interrupt.

If the MSC supports signaling the MPAM error interrupt through a *Message signaled interrupt* (MSI), the interrupt must be edge-triggered.

See also [Chapter 10 Errors in MSCs](#).

6.8.1.2 Edge-triggered interrupts

When using edge-triggered interrupts, the interrupt edge is generated when `MPAMF_ESR.ERRCODE` is written due to an error.

An edge-triggered interrupt is not generated when software writes to `MPAMF_ESR`.

An interrupt service routine does not need to clear an edge-triggered interrupt.

If the MSC supports signaling the MPAM error interrupt through an MSI, the interrupt must be edge-triggered.

See [Chapter 10 Errors in MSCs](#) for other reasons for an interrupt service routine to clear `MPAMF_ESR`.

6.8.1.3 Support for MSI writes to signal error interrupts

Message signaled interrupts (MSIs) are signaled using a memory write that is usually directed at an interrupt translation service.

The support for error MSIs is identified by the `MPAMF_IDR.{HAS_ERR_MSI, HAS_ESR}` fields.

The registers that contain the error MSI write configuration are:

- `MPAMF_ERR_MSI_ADDR_L`.
- `MPAMF_ERR_MSI_ADDR_H`.
- `MPAMF_ERR_MSI_ATTR`.
- `MPAMF_ERR_MSI_DATA`.
- `MPAMF_ERR_MSI_MPAM`.

Instances of these MSI configuration registers exist in each of the Secure physical address space and the Non-secure physical address space. The set of these registers in an address space configures the error MSI write for errors from the `MPAMCFG_*` or `MPAMF_*` registers in that address space.

Errors can also be raised by errors in requests. Errors in requests which have the PARTID space selected by `MPAM_NS` of 0 are signaled as Secure errors using the MSI write information from the `MPAMF_ERR_MSI_*` registers in the Secure address space. Errors in requests which have the PARTID space selected by `MPAM_NS` of 1 are signaled as Non-secure errors using the MSI write information from the `MPAMF_ERR_MSI_*` registers in the Non-secure space.

6.8.2 MPAM overflow interrupt

A monitor can overflow, especially if it is a type of monitor that accumulates counts. If it is possible for a type of monitor to overflow, there are bits in `MSMON_CFG_*_CTL` to control the behavior on overflow (*Overflow status bit*).

Support of an overflow interrupt is optional in an MSC. If the MSC has monitors that can overflow, Arm recommends that the MPAM overflow interrupt be implemented.

When an MPAM monitor instance overflows, it sets the `OFLOW_STATUS` flag in the monitor instance's control register. If the `OFLOW_STATUS` flag was previously 0 and `OFLOW_INTR` bit is 1, an overflow interrupt is signaled if the MSC implements overflow interrupts.

If an MSC supports both Secure and Non-secure address spaces, `MSMON_CFG_*_CTL` registers and `MSMON_MBWU` and `MSMON_CSU` registers that are implemented have Secure and Non-secure instances. Secure instances of `MSMON_CFG_*_CTL.OFLOW_INTR` control whether a Secure MPAM overflow interrupt is

generated when the corresponding Secure counter instance overflows. Non-secure instances of `MSMON_CFG_*_CTL.OFLOW_INTR` control whether a Non-secure MPAM overflow interrupt is generated when the corresponding Non-secure counter instance overflows.

- The mechanism by which an interrupt request from an MSC resource monitor generates an FIQ or IRQ exception is IMPLEMENTATION DEFINED.
- Arm recommends that an MSC implements two MPAM overflow interrupt signals, one for the Secure MPAM overflow interrupt and another for the Non-secure MPAM overflow interrupt.
- Arm recommends that MPAM overflow interrupt requests:
 - Translate into an `MPAM_OF_IRQ` signal, so that they are observable to external devices.
 - If the MSC is integrated into a PE, connect to inputs on an IMPLEMENTATION DEFINED generic interrupt controller as a Private Peripheral Interrupt (PPI) or a Locality-specific Peripheral Interrupt (LPI) for that PE. See the *Arm Generic Interrupt Controller Architecture Specification, GIC architecture version 3.0 and version 4.0* for information about PPIs, LPIs and SPIs.
 - If the MSC is not integrated into a PE, connect to inputs on an IMPLEMENTATION DEFINED generic interrupt controller as a System Peripheral Interrupt (SPI) or Local Peripheral Interrupt (LPI).

The interrupt is reset by writing 0 to the `OFLOW_STATUS` field of all overflowed monitor instances `MSMON_CFG_*_CTL` register.

If the MSC supports signaling monitor overflow interrupts through an MSI, the MPAM monitor overflow interrupt must be edge-triggered.

6.8.2.1 Support for MSI writes to signal overflow interrupts

MSIs are signaled using a memory write that is usually directed at an interrupt translation service.

The support for the monitor overflow interrupt is identified by the `MPAMF_MSMON_IDR.{HAS_OFLW_INTR, HAS_OFLW_MSI}` fields.

The registers that contain the error MSI write configuration are:

- `MSMON_OFLOW_MSI_ADDR_L`
- `MSMON_OFLOW_MSI_ADDR_H`.
- `MSMON_OFLOW_MSI_ATTR`.
- `MSMON_OFLOW_MSI_DATA`.
- `MSMON_OFLOW_MSI_MPAM`.

Instances of these MSI configuration registers exist in each of the Secure physical address space and the Non-secure physical address space. The set of these registers in an address space configures the overflow MSI write from overflow events of monitors accessible in that address space.

6.8.2.2 Monitor overflow status register

The optional `MSMON_OFLOW_SR` register gives a summary of the overflow status flags (`OFLOW_STATUS` and `OFLOW_STATUS_L`) for each RIS and for each monitor type.

This register contains a flag bit per RIS value. Each flag is 0 if all of the `OFLOW_STATUS` and `OFLOW_STATUS_L` bits of all monitor types and all instances of each type for the resource instance are 0. Each flag is 1 if any of the overflow status bits for any monitor instance of any type for the resource instance are 1.

The register also contains a flag bit for each monitor type. A monitor type flag is 1 if any monitor instance of the type for the resource instance has the `OFLOW_STATUS` or `OFLOW_STATUS_L` bit as 1.

`MSMON_OFLOW_SR` is read-only. The flags are reset when the `OFLOW_STATUS` and `OFLOW_STATUS_L` bits monitored by that flag have all be reset to zero.

The presence of `MSMON_OFLOW_SR` is indicated by `MPAMF_MSMON_IDR.HAS_OFLOW_SR == 1`.

6.8.2.3 Monitor type overflow status bitmap registers

In an implementation that has many monitor instances of a monitor type, the number of monitor instances to scan for overflows is large even after consulting `MSMON_OFLOW_SR` to eliminate most of the RIS and monitor types. To probe one monitor instance requires that the monitor overflow interrupt service routine set `MSMON_CFG_MON_SEL` to a monitor instance, read `MSMON_CFG_<type>_CTL` and check one or two bits in that register to see if the `OFLOW_STATUS` or `OFLOW_STATUS_L` bit is set.

To assist the scanning of many monitor instances, optional overflow status bitmap registers for a monitor type are available for implementation. These overflow status bitmaps can greatly accelerate the scanning.

Each MPAM monitor type can have an optional overflow status register that shows the overflow status flags in a bitmap of 32 monitor instances. The monitor instances shown are selected in `MSMON_CFG_MON_SEL` where the RIS field selects the resource instance and the `MON_SEL` field AND `0xFFE0` selects the lowest of the contiguous 32 monitor instances reported in the bitmap.

For the CSU monitor type, the CSU overflow status register is `MSMON_CSU_OFSR`. The presence of this register is discoverable in `MPAMF_CSUMON_IDR.HAS_OFSR`.

———— Note ————

In most implementations, CSU monitor instances will not be able to overflow as the maximum value in `MSMON_CSU` is known at design time and will fit within the architectural maximum of `MSMON_CSU`. In such an implementation, there will be no CSU monitor instance overflows and `MSMON_CSU_OFSR` has no value.

For the MBWU monitor type, the MBWU overflow status register is `MSMON_MBWU_OFSR`. The presence of this register is discoverable in `MPAMF_MBWUMON_IDR.HAS_OFSR`.

6.9 MSC support of MPAM for RME

An RME system supports 4 physical address spaces. MPAM for RME supports the 4 address spaces and 4 PARTID spaces. The MPAM system environment of an RME system is described in *The MPAM for RME system*.

Requirements and definitions are:

- An MSC that supports a PARTID space must support the associated Physical Address Space that is needed for accessing the control settings configurations of the PARTIDs in that PARTID space.
- An MSC that supports the 4 physical address spaces and 4 PARTID spaces is defined as a *four space MPAM MSC*.
- An MSC that supports either 4 or 2 physical address spaces and 2 PARTID spaces is defined as a *two space MPAM MSC*.
- Non-MPAM components support either 1, 2 or 4 address spaces but do not support MPAM at all. Non-MPAM devices have no regulated resources and must not have MPAM devices downstream. See *Non-MPAM components*.
- Other combinations of physical address space support and PARTID space support are not permitted.

4 PARTID spaces must be supported in the levels of interconnect that connect RME PEs, but some MSCs might support MPAM with support for only 2 PARTID spaces. See *MPAM for RME propagation of MPAM_SP with requests*.

The MPAM PARTID space in a request and the physical address space accessed by the request are independent in the request. The associations of physical address space and PARTID space are part of the request generation process at a Requester. An MSC must not assume any association between the PARTID space of a request and the physical address space of the request.

A four space MSC is permitted to use either a Non-secure MPAM error interrupt or a Secure MPAM error interrupt for reporting an error associated with the Root or Realm PARTID spaces.

6.9.1 Four-space MSC

An MSC that fully supports RME and MPAM must have 4 PARTID spaces and 4 physical address spaces.

In an MSC that supports 4 PARTID spaces and 4 physical address spaces, the `MPAMF_IDR.SP4` bit must be 1 when read from any address space and, if RIS is supported, with any `MPAMCFG_PART_SEL.RIS` value.

`MPAMF_BASE_s`, `MPAMF_BASE_ns`, `MPAMF_BASE_rt`, `MPAMF_BASE_rl` must all be defined in the firmware table description of the MSC.

The MPAM memory-mapped registers in each address space are at the offsets from the MPAM Feature Page Base address in that address space. [Table 6-2](#) shows the relationship of address space, the MPAM feature page base address symbol and the contents of that MPAM feature page.

Table 6-2 Relationship of address space, MPAM feature page base address symbol and a description of the contents of that MPAM feature page

Address Space	MPAM Feature Page Base	Description
Non-Secure	<code>MPAMF_BASE_ns</code>	MPAM MSC registers in the Non-secure address space describe and access controls and monitors for Non-secure PARTID space.
Secure	<code>MPAMF_BASE_s</code>	MPAM MSC registers in the Secure address space describe and access controls and monitors for Secure PARTIDs.
Realm	<code>MPAMF_BASE_rl</code>	MPAM MSC registers in the Realm address space describe and access controls and monitors for the Realm PARTID space.
Root	<code>MPAMF_BASE_rt</code>	MPAM MSC registers in the Root address space describe and access controls and monitors for the Root PARTID space.

The offsets of MPAM memory-mapped registers from the MPAM Feature Page base address are the same for each MPAM Feature page and in each address space. See [Table 9-1](#) for all MPAM MSC registers. Added fields and accessors for the two physical address spaces for RME are described in this chapter. See [Chapter 9 Memory-mapped Registers](#) for Memory-mapped registers from the MPAMF_BASE_* for that address space.

See [Minimum required MPAM memory-mapped registers](#) for the required minimum set of MPAM registers accessible from the MPAM Feature Page in any address space. In each address space the MPAM features of the MSC in that address space are described by decoding the fields in MPAMF_IDR. This indicates that additional ID registers are present and further describe the features. MPAM has no requirement that the resource controls and monitors in one address space are the same as those described in another address space.

Instances of the MPAMCFG_* registers must exist in each of the 4 address spaces where MPAMF_*IDR.HAS_* is 1 for a feature that uses those registers.

There must be an instance of MPAMCFG_PART_SEL in each of the 4 address spaces unless there are no resource controls or resource instances in the PARTID space whose control registers are accessed through that physical address space.

Instances of the MSMON_* registers must exist in each address space where the ID registers indicate that the monitor exists.

There must be an instance of MSMON_CFG_MON_SEL in each of the 4 address spaces that contain any monitor registers.

MPAMF_ESR and MPAMF_ECR must exist in each address space in each of the 4 address spaces where MPAMF_IDR.HAS_ESR is 1.

Chapter 7

Resource Partitioning Controls

This chapter contains the following sections:

- *Introduction.*
- *MPAM partitionable resources.*
- *Standard partitioning control interfaces.*
- *Vendor or implementation-specific partitioning control interfaces.*
- *Measurements for controlling resource usage.*
- *PARTID narrowing.*
- *System reset of MPAM controls in MSCs.*
- *About the fixed-point fractional format.*

7.1 Introduction

This introduction to memory-system partitioning is *informative*. Other sections are *normative* unless marked as *informative*.

Software assigns VMs and applications to a partition. The hypervisor can assign VMs to partitions, and operating systems can assign applications to partitions. This specification does not address how such assignments are made by software.

A memory-system partition is associated with a software environment on a PE by loading an MPAMn_ELx register with PARTID_I and PARTID_D. An EL2 hypervisor loads MPAM1_EL1 with the partition IDs when context-switching between VMs. An EL1 operating system loads MPAM0_EL1 with the partition IDs when context-switching between applications. The PARTIDs loaded into fields of MPAMn_ELx for instruction and data accesses are used for requests when running software at ELn. The PARTID on memory-system requests connects the software environment to the resource partitioning controls in the MSCs that handle the requests.




Figure 7-1 Partitioning, VMs, and OS processes

The PARTID of a request controls uses of each MSC’s performance resources. An MSC receives a PARTID with each request. The PARTID may be used within the component to select resource controls for the component’s resource allocation and utilization behavior.

All memory-system requests with a given PARTID share the resource control settings for that partition.

Because a PARTID is communicated to shared MSCs and interpreted there, PARTIDs should be managed and allocated on a system-wide basis.

Resource partitioning controls might be standard or implementation specific.

Standard control interfaces are architected, but optional. Therefore, an MSC that does not require a standard control interface does not need to implement it. Most MSCs implement few of the standard control interfaces.

An implementation-specific resource control can use a PARTID for unique facilities that either control resources not envisioned by the standard controls or that implement unique control methods that cannot be mapped onto the standard control interfaces.

7.2 MPAM partitionable resources

An MSC contains resources that affect the performance of the memory system. For such a resource to be partitionable:

- The component must support MPAM at its upstream interface.
- The component must have one or more MPAM resource controls for that resource.

A partitionable resource may be partially allocated to a partition according to the programming of the MPAM resource control or controls for that resource.

If the implementation supports the RIS MPAM feature, the MSC may have two or more partitionable resources differentiated by the value of `MPAMCFG_PART_SEL.RIS`. For more information see [Resource instance selection](#).

7.3 Standard partitioning control interfaces

The MPAM architecture defines standard partitioning control interfaces. This enables binary distribution of operating systems supporting MPAM.

The MPAM architecture defines the following standard types of control interfaces for memory-system resources:

- Cache-portion partitioning.
- Cache maximum-capacity partitioning.
- Cache maximum associativity partitioning.
- Memory-bandwidth portion partitioning.
- Memory-bandwidth minimum and maximum partitioning.
- Memory-bandwidth proportional-stride partitioning.
- Priority partitioning.

Each of these standard control interfaces is optional at each MSC. An MSC may implement several controls or none. Some controls only make sense for certain types of MSCs, or for certain implementations of an MSC. Others may be possible but too costly for the system's target market.

Cache-portion partitioning and memory-bandwidth portion partitioning follow the generic portion-control interface described in *Portion resource controls*. Cache maximum-capacity partitioning follows the generic maximum-usage control interface described in *Maximum-usage resource controls*.

The presence of each standard control is indicated by a bit in `MPAMF_IDR`, or in a resource-specific memory-mapped ID register. See *Memory-mapped ID register description*.

The control settings storage is accessed through the combination of several access indices:

- The address space used to access the Secure or Non-secure MSC register. Controls for PARTIDs in:
 - The Secure PARTID space are accessed through registers in the Secure address space
 - The Non-secure PARTID space are accessed through registers in the Non-secure address space.
- The MSC that contains the control. This is represented as the base address of the MPAM feature page in the address space. These are represented here as:
 - `MPAMF_BASE_s` in the Secure address space.
 - `MPAMF_BASE_ns` in the Non-secure address space.
- If `MPAMF_IDR.HAS_RIS` is 1, `MPAMCFG_PART_SEL.RIS`. This field selects a resource to access.
- `MPAMCFG_PART_SEL.PARTID`. This field selects the PARTID from:
 - The PARTID space.
 - The resource instance to be configured.
- The control settings register. When accessed, this register selects which control is being configured for:
 - The PARTID.
 - The PARTID space.
 - The resource instance.

For example, to access the memory bandwidth maximum configuration settings for Secure PARTID 15 on resource instance 2 of an MSC that implements RIS:

1. Secure PARTID 15 must be stored in `MPAMCFG_PART_SEL.PARTID` at the address `MPAMF_BASE_s + 0x0100` and, due to RIS being implemented, the `RIS` field of that address must be set to 2 to ensure access to the correct resource instance.
2. Once the store has completed, the new maximum fraction of memory bandwidth for Secure PARTID 15 of resource instance 2 must be stored into the `MPAMCFG_MBW_MAX_s` register of this MSC, found at `MPAMF_BASE_s + 0x0208`.

Software must ensure mutual exclusion for access to `MPAMCFG_*` registers of each MSC.

7.3.1 Disabling a PARTID

A PARTID can be enabled or disabled with a single store to the MSC. The enabled status of a PARTID within a PARTID space is global to the MSC. It applies to all the other partitioning controls in that MSC.

This functionality allows for the rapid reclaim of resources used by a PARTID if the software using the PARTID is exited or at equipment-wide mode changes.

A disabled PARTID behaves as if it is programmed to allocate no resource or allocate resource only with the lowest allocation priority.

A PARTID is disabled by `MPAMCFG_DIS` with the PARTID and an NFU bit. The NFU bit declares that software has no future use of that PARTID. If `MPAMCFG_DIS.NFU` is written as 0, hardware must preserve the settings of `MPAMCFG_DIS.PARTID` and those settings must immediately be used once the PARTID is enabled at a later time. If `MPAMCFG_DIS.NFU` is written as 1, hardware is permitted to either preserve the settings or to discard them and so requiring new settings to be configured before the PARTID is re-enabled.

The PARTID disabled or enabled is in the PARTID space that corresponds to the `MPAMF_BASE` page instance used to access the `MPAMCFG_DIS` or `MPAMCFG_EN` register.

The settings for a disabled PARTID can be accessed to read or write the control settings. See `MPAMCFG_PART_SEL` and *Resource controls*.

If *Resource instance selection* is enabled in the MSC, The `MPAMCFG_DIS`, `MPAMCFG_EN` and registers affect the behavior of all resource instances in the PARTID.

Instances of `MPAMCFG_EN`, `MPAMCFG_DIS` and `MPAMCFG_EN_FLAGS` exist at fixed offsets in each MPAM feature page. There is an MPAM feature page in each physical address space. These feature pages access the control settings of the PARTID in the PARTID space associated with that address space by using the `MPAMCFG_PART_SEL` register instance in the same physical address space.

To assist in enabling or disabling many PARTIDs, `MPAMCFG_EN_FLAGS` accesses the enable flags of 32 PARTIDs at once for reading or writing as a bit vector. The block of PARTID enable flags are from `MPAMCFG_PART_SELPARTID_SEL & 0xFFE0` in `MPAMCFG_EN_FLAGS[0]` for 32 PARTIDs to `MPAMCFG_PART_SELPARTID_SEL & 0xFFE0 + 31` in `MPAMCFG_EN_FLAGS[31]`.

When a PARTID is disabled by writing to `MPAMCFG_EN_FLAGS`, hardware must preserve the settings for that PARTID.

A PARTID is re-enabled by storing `MPAMCFG_EN` with the PARTID to enable.

The enabled status of a PARTID controls the behavior of the request PARTID, even when *PARTID narrowing* is implemented.

7.3.1.1 Enabled and disabled behavior of resource controls

The behavior of an enabled control is the normal behavior of that control with the programmed control settings.

The behavior of a disabled control depends on the type of the control:

- A fractional control is the same as if the control were set to zero.

- A portion-based partitioning control is as if the control were set to allocate no portions.
- In a cache, a disabled PARTID has the second lowest allocation priority. This is higher than the allocation priority of an empty cache line, meaning it is likely to be replaced.
- For resource controls, only the control for PARTID of 0 must be reset. PARTID 0 must be reset to enabled 0b1. Firmware must initialize all PARTIDs in all PARTID spaces to be enabled at system boot before transferring to software that might not support this feature.
- When *PARTID narrowing* is implemented, the resource control settings could be shared by multiple request PARTIDs. Therefore, disabling a PARTID cannot make the resource control for a shared internal PARTID act as if it has no resource. Instead it must act as if the request PARTID is temporarily mapped to an internal PARTID that has no resources assigned. This means that `MPAMCFG_DIS.NFU` must not be implemented in an MSC that has *PARTID narrowing*.

7.3.2 Cache-portion partitioning

A portion is a uniquely identifiable part of a resource. It is of fixed size or capacity and all portions of a resource are the same size. A particular resource has a constant number of portions. Every partition that is given access to a portion *n* shares access to portion *n*.

The storage portions of caches can be partitioned. Allocating portions of a cache to a partition permits requests attributed to that partition to allocate within those portions of the cache.

When a request to a cache requires a cache line to be installed in the cache, the PARTID of that request determines which portions of the cache the request may allocate to install the line.

Cache-portion partitioning uses the generic portion-partitioning interface described in *Portion resource controls*.

7.3.2.1 Cache-portion bit map

A cache-portion bitmap (CPBM) controls the cache-storage portion allocation for a partition. Each bit of a CPBM controls whether the partition is permitted to allocate a particular capacity portion of the cache. The number of capacity portions available in a cache is an IMPLEMENTATION DEFINED parameter that is discoverable in `MPAMF_CPOR_IDR` for the cache. The width of the CPBM field is equal to the number of capacity portions available in the cache.

For example, assume a cache has a 1 MB total capacity in 32 portions. Each portion has a capacity of $1\text{ MB} / 32 = 32\text{ KB}$. A partition has 4 portions allocated (only 4 bits in the CPBM are 1's). So, this partition can only allocate into these particular 4 portions, allowing up to 128 KB, or 1/8th of the cache's total capacity.

CPBM is an instance of the generic portion bitmap (PBM) described in *Portion resource controls*.

7.3.2.2 Over-allocation of capacity portions

Storage capacity portions cannot be over-allocated. This is true because the CPBM contains bits that control allocations in the implementation-dependent number of allocable capacity portions of the cache.

7.3.2.3 Changing CPBM for a partition

Software may change the CPBM during system operation. This does not disrupt normal system operation because the CPBM only affects new allocations and does not reallocate previously allocated cache storage.

If a cache line was allocated under a previous CPBM to a portion that is not set in the new CPBM, the partition is using more of the cache capacity than it is entitled to under the new CPBM:

- If lines previously allocated in a portion that is not in the new CPBM are not accessed again, they will eventually be reallocated to a partition that has its CPBM bit set for that portion of the capacity. So, these will represent a temporary mis-allocation of capacity.

- If however, a line that is present in the cache in a portion that is not in the new CPBM continues to be accessed, this can lead to a long-term mis-allocation of capacity. The line's location optionally might be updated, see *Write hits that update the resource partition of a cache line may move that line to a different portion*.

7.3.2.4 Using cache-portion partitioning with cache maximum-capacity partitioning

When cache-portion partitioning is used with cache maximum-capacity partitioning, both controls are effective as described in *Using cache maximum-capacity partitioning with cache-portion partitioning*.

7.3.3 Cache minimum-capacity resource control (CMIN)

Cache minimum-capacity resource control is a memory bandwidth minimum control that gives priority to a portion of the cache capacity that a PARTID can use. When implemented this control is used in addition to cache maximum capacity.

The `MPAMCFG_CMIN` register has the CMIN control setting in 16-bit fixed-point fraction format. A PARTID that currently occupies less than CMIN fraction of the cache capacity has an elevated allocation priority.

The implemented bit width of the CMIN field is described in `MPAMF_CCAP_IDR.CMAX_WD`. Implemented bits are always bit [15:16-CMAX_WD]. If both `MPAMCFG_CMAX` and `MPAMCFG_CMIN` are implemented, the implemented most significant fractional bits of both registers are the same, and that width is `CMAX_WD`.

7.3.3.1 CMIN and Cache allocation priorities

The partitioning control contributes one level in a cache line allocation priority scheme, as shown in [Table 7-1](#)

Table 7-1 Cache allocation priorities

Cache line occupant	Priority	Description
Unallocated	0 - lowest	Always allocatable
Disabled PARTID	1	Line occupied by disabled PARTID
PARTID with capacity over CMAX	2	Line occupied by PARTID using > CMAX
PARTID between CMAX and CMIN	3	Line occupied by PARTID in midrange
PARTID under CMIN	4	High priority

The request PARTID is evaluated against these criteria and assigned a priority. A request cannot have priority of 0.

The priority values of each line in the set are compared and if there is at least one current cache occupant with lower priority, the request replaces one lowest priority of these.

If there are no lower priority lines, but one or more lines of the same priority, the request can replace one of those according to the following exceptions:

- If the request PARTID is currently using > CMAX of the capacity and has `MPAMCFG_CMAX.SOFTLIM == 0`, the request may only replace a line currently occupied by the same PARTID.
- If the *Cache associativity partitioning (CASSOC)* control is implemented, this may limit the request PARTID to not allocate any new capacity in the associativity unit (cache set) but only to replace one of the request PARTID's previously allocated cache lines.
- Otherwise, the replacement target is chosen from among the lines at the same priority or lower by some other mechanism, presumably the cache replacement algorithm.

If there are no current occupants with the same or lower priority than the request PARTID, then the request does not allocate in the cache.

———— **Note** ————

A request by a PARTID always ties with lines currently occupied by the same PARTID as this priority is determined solely from the PARTID.

7.3.3.2 Layered filtering of allocation choices

The choice of allocation candidates involves the following steps. They must be performed in order:

1. All implemented MPAM cache partitioning controls are applied:
 - If `MPAMCFG_CPBM<n>` is implemented, only candidates that have the `MPAMCFG_CPBM<n>` bit set are included.
 - If `MPAMCFG_CMAX` is implemented and the PARTID currently occupies greater than CMAX fraction of the cache capacity, the choices are limited to:
 - If `MPAMCFG_CMAX.SOFTLIM` is 0, lines currently occupied by the same PARTID are treated as the set of candidates.
 - If `MPAMCFG_CMAX.SOFTLIM` is 1 specifying Unallocated lines, those occupied by a disabled PARTID and those occupied by the same PARTID are treated as the set of candidates.
 - If *Cache associativity partitioning (CASSOC)* is implemented and the request PARTID currently occupies more than the CASSOC fraction of the associativity in the unit of associativity that the request addresses, the choices are limited to only those lines already occupied by the request PARTID.
2. The CMIN priority of the request is compared to each of the remaining lines and those candidates that are occupied by a PARTID of higher priority than the request are removed as candidates.
3. If no replacement candidates remain, no line is allocated. If candidates exist, the implementation's cache replacement algorithm, for example LRU, is used to select between the remaining candidates.

7.3.4 Cache associativity partitioning (CASSOC)

Cache associativity partitioning gives direct control over the amount of associativity that the PARTID may use within any unit of associativity. In a set associative cache, the CASSOC control sets the maximum fraction of the ways that a PARTID could allocate within any cache set. In a fully associate cache, it would set the fraction of the entries that the PARTID could use.

The cache maximum associativity control setting for a PARTID is accessed through the `MPAMCFG_CASSOC.CASSOC` field. This field is encoded in the fixed-point fraction format.

Maximum cache associativity usage in fixed-point fraction format by the partition selected by `MPAMCFG_PART_SEL`.

CASSOC sets the maximum fraction of the cache associativity that the PARTID is permitted to allocate. CASSOC regulates the associativity in each associativity grouping of the cache. In a set associative cache, CASSOC applies to the fraction of the ways in each set.

The implemented width of the fixed-point fraction is defined by `MPAMF_CCAP_IDR.CASSOC_WD`.

Unimplemented bits within the `MPAMCFG_CASSOC.CASSOC` field are RAZ/WI. The implemented bits of the CASSOC field are always the most-significant bits of the field.

The fixed-point fraction CASSOC is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the cache that can be represented in an implementation with w implemented bits is $1 - 1/(2^w)$.

As with the other resource controls, the control for PARTID must reset to all 1s, the maximum fraction.

An instance of `MPAMCFG_CASSOC` exists at a fixed offset in each MPAM feature page. There is an MPAM feature page in each physical address space. Accesses the CASSOC control setting for the PARTID in the PARTID space associated with that address space using the `MPAMCFG_PART_SEL` register instance in that same physical address space.

7.3.5 Cache maximum-capacity partitioning

A limit may be set on the storage capacity of a cache that a memory-system partition may use. Setting a maximum cache capacity to a partition permits requests attributed to that partition to allocate up to that maximum cache capacity. Attempts to allocate beyond that capacity must limit a partition's capacity usage.

Techniques for limiting cache usage by a new request when a partition's capacity usage is at or above its maximum include:

- Do not allocate for the new request.
- Replace some data from that partition with data from the new request.
- Evict some data from that partition from the cache before allocating for the new request.
- Defer the required deallocation until a more convenient time.

Cache lookups are not affected by partitioning. A cache lookup must find a valid cache line even if that line was allocated with a different PARTID.

Cache maximum-capacity partitioning follows the description of the generic maximum-usage resource control interface described in *Maximum-usage resource controls*.

7.3.5.1 Cache maximum-capacity control setting

The cache maximum-capacity control setting is programmed by storing a capacity limit into the MSC's cache maximum capacity control interface, `MPAMCFG_CMAX`.

This setting is a hard limit on the cache capacity. To set a soft limit, see *Cache maximum-capacity control soft limit*.

The cache maximum-capacity limit is a fraction of the cache's total capacity. The format of the limit value is a fixed-point fraction, as described in *About the fixed-point fractional format*.

For example, to allocate 30% of a 256 KB cache to a partition:

- In the fixed-point fractional format, 1.0 is represented as $2^{16} - 1$, or in hex as `0xFFFF`. The subtraction makes 1.0 within the range of the representation.
- So, the representation of 30% would be $1.0 * 0.30$, which in hex is `0xFFFF * (decimal) 0.30`, or `0x4CCC`.
— Similarly, 25% would be `0x3FFF`; 14% would be `0x23D6`; 3% would be `0x07AE`; and 3.25% would be `0x0851`.
- If you have a cache with 256 KB of capacity, and the resource control setting for a PARTID is set to `0x4CCC` to represent 30%, that partition is permitted to use 30% of the cache, or about 76.75 KB of capacity.
- Since most, but not all, Arm caches have 64-byte lines, a 256 KB cache has 4096 of these 64-byte lines, and 30% of those lines is 1228 or 76.75 KB.

The fixed-point fractional format permits an implementation to leave bits to the right as unimplemented, meaning that the value would be truncated to the implemented bits, causing some of the right-most bits to be zeros:

- As an example, the 3% value previously mentioned is `0x07AE`. If only 8 bits of fraction are implemented, when software stores `0x07AE` into a resource control setting, the value is shortened to the most significant bits and stored as `0x07--`.
- When using the resource control setting, the unimplemented bits would be read as zeros.

The actual value of the setting is therefore an interval from the value of the control setting up to the value of the control setting plus one in the right-most implemented bit.

- In the case of the 3% value previously mentioned, that interval is from 0x07 (2.734%) to 0x08 (3.125%).
- An implementation is permitted to regulate the resource to any point within this interval.

7.3.5.2 Cache maximum-capacity control soft limit

When `MPAMF_CCAP_IDR.HAS_CMAX_SOFTLIM` is 1, `MPAMCFG_CMAX` implements the `SOFTLIM` field.

When `MPAMCFG_CMAX.SOFTLIM` is 0, the cache maximum capacity control sets a hard limit that prevents the `PARTID` from allocating more than the maximum fraction of the cache capacity.

When a `PARTID`'s request requires allocation in the cache but the `PARTID`'s cache capacity usage is above the fraction of the cache capacity set by its `MPAMCFG_CMAX` register, its capacity use cannot be increased even if there is unused capacity.

Setting the `CMAX` control for a `PARTID` to be set to a soft limit allows unused capacity to be temporarily used. The control bit `MPAMCFG_CMAX.SOFTLIM` controls this behavior. When `SOFTLIM` is 1, a `PARTID` that is currently using more than its `MPAMCFG_CMAX.CMAX` capacity is permitted to allocate more capacity from an Unallocated line or a line that is in use by a disabled `PARTID`.

`MPAMCFG_CMAX.SOFTLIM` for `PARTID 0` must reset to 0, hard limit behavior. For compatibility with software that does not support `SOFTLIM`, the firmware must reset `SOFTLIM` to 0 for all `PARTID`s in all `PARTID` spaces before transferring to software that might not support `SOFTLIM`.

7.3.5.3 Using cache maximum-capacity partitioning with cache-portion partitioning

When cache-portion partitioning is used with cache maximum-capacity partitioning, both controls are effective. Cache-portion partitioning controls which portions of the capacity may be allocated to this partition. Cache maximum-capacity partitioning limits the amount to less than or equal to a cache-capacity limit control setting.

For example, assume several portions of the capacity are shared by several partitions. Any such partition can allocate within the shared portions. To keep one of the partitions from using too much of the shared allocation, the maximum-capacity controls for the partitions can each be set to less than the capacity of the portions to which they may allocate. If each partition is given 50% of the capacity of the shared portions, then no one partition can use more than 50% of the shared cache portions.

Here is an example of a cache with 1 MB total capacity in 32 portions. Each partition has 4 portions for shared allocation. To allow a partition to use no more than 50% of its shared allocation, you would set the cache maximum-capacity limit for this partition as follows:

1. Portions divide the capacity of the cache into distinct parts of the same size. So, for a 1 MB cache divided into 32 portions, each portion has $1 \text{ MB} / 32 = 32 \text{ KB}$:
 - a. In portion partitioning, it is not possible to allocate anything other than an integral number of portions to a `PARTID`.
 - b. A cache portion may be exclusively allocated to a `PARTID` or it may be shared by 2 or more `PARTID`s.
 - c. A `PARTID` that has 4 portions allocated to it is permitted to use $32 \text{ KB} * 4 = 128 \text{ KB}$.
2. The combined behavior of cache-portion partitioning and cache maximum-capacity control has both controls:
 - a. To allow a `PARTID` to use only 50% of the storage in the portions allocated to it, the cache maximum-capacity control is used.
 - b. Compute the fraction of the cache that is 50% of the storage in the portions allocated. In this case, it is $64 \text{ KB} / 1 \text{ MB} = 1/16$ or 6.25%, which is 0x0FFF in the fixed-point fractional representation.

- c. The combined behavior only permits the PARTID to allocate storage in the 4 portions it may use according to the cache-portion control, but its use of storage is also limited to 50% of the storage of those portions.

7.3.5.4 Over-allocation of capacity

Cache capacity can be over-allocated because the sum of the cache-capacity control parameters may exceed 100% of the cache size. This may be acceptable. The cache-capacity control does not provide a minimum cache capacity guarantee, only a maximum guarantee. The data of inactive partitions may be evicted from the cache due to the activity of other partitions.

7.3.6 Memory-bandwidth portion partitioning

An MSC's downstream bandwidth may be divided into portions, and those portions may be allocated to partitions.

Memory-bandwidth portion partitioning follows the generic portion-control interface described in *Portion resource controls*, in which a portion is a quantum of bandwidth. A Time-Division Multiplexing (TDM) scheme that allocates traffic to time slots is an example of a bandwidth allocation system that has portions.

The BandWidth Portion Bit Map (BWPBM) is the Portion Bit Map (PBM) for bandwidth.

7.3.7 Memory-bandwidth minimum and maximum partitioning

An MSC's downstream bandwidth may be partitioned by bandwidth usage. There are two bandwidth-usage control schemes. An MSC can optionally implement each of them:

- Minimum bandwidth to which the PARTID has claim, even in the presence of contention.
- Maximum bandwidth limit available to the PARTID, in the presence of contention.

The minimum and maximum bandwidth partitioning schemes rely on tracking bandwidth usage by PARTIDs. Because bandwidth is measured in bytes per second, bandwidth measurements have a dependence on time. That dependence is captured in this specification as the accounting window or accounting period. See *Memory-bandwidth allocation accounting window width*

Without contention, the bandwidth may be strictly limited to the maximum or permitted to use more than the maximum, since no other partition's traffic is claiming that bandwidth.

Any combination of these control schemes may be used simultaneously in an MSC that supports them.

Each control scheme is described below.

7.3.7.1 Minimum-bandwidth limit partitioning

The minimum-bandwidth control scheme regulates the bandwidth used by a PARTID's requests:

- If the bandwidth usage by the PARTID of the request, as tracked during the accounting period, is currently less than the partition's minimum, its requests are preferentially selected to use downstream bandwidth.
- If the bandwidth usage by the PARTID of the request, as tracked during the accounting period, is currently greater than or equal to the PARTID's minimum, its requests compete with other requests as described under *Maximum-bandwidth limit partitioning*, if implemented. If maximum-bandwidth limit partitioning is not implemented, requests with PARTID that have current bandwidth usage greater than that PARTID's minimum-bandwidth limit compete with all requests and do not receive preferential treatment under the minimum-bandwidth limit.

A PARTID's requests below its minimum bandwidth are therefore most likely to be scheduled to use downstream bandwidth.

Bandwidth that is not used by a partition during an accounting window does not accumulate.

The control parameter is a fixed-point fraction of the available bandwidth. For more information, see *About the fixed-point fractional format*.

7.3.7.2 Maximum-bandwidth limit partitioning

The maximum-bandwidth limit control scheme regulates the bandwidth used by a PARTID's requests:

- If the bandwidth usage by the PARTID as tracked during the accounting period is currently less than the PARTID's maximum bandwidth but greater than or equal to its minimum bandwidth, if implemented, its requests are selected to use bandwidth when there are no competing minimum bandwidth requests to service. Requests for PARTIDs that are above their minimum-bandwidth limits but less than their maximum-bandwidth limits compete with each other to use bandwidth.
- If the bandwidth usage by the PARTID of the request is greater than or equal to the PARTID's maximum bandwidth and the HARDLIM bit is not set, the request competes with other such requests to use bandwidth when there are no competing requests to service for PARTIDs currently below their minimum bandwidth or maximum bandwidth.
- If the bandwidth usage by the PARTID of the request is greater than or equal to the PARTID's maximum bandwidth and the Hard Limit (HARDLIM) bit is set, the requests are saved until the PARTID's bandwidth usage drops below its maximum bandwidth control setting.

If the HARDLIM bit is set, the partition is prevented from using more bandwidth if the current bandwidth usage is over the maximum bandwidth limit. As the accounting window advances, the current bandwidth usage resets to zero or otherwise decays, permitting the partition to again use bandwidth.

Bandwidth that is not used by a partition during an accounting window does not accumulate.

The control parameter is a fixed-point fraction of the available bandwidth. For more information, see [About the fixed-point fractional format](#).

7.3.7.3 Using minimum-bandwidth limit with maximum-bandwidth limit controls

If both minimum-bandwidth limit and maximum-bandwidth limit are implemented, [Table 7-2](#) shows the preference of requests.

Table 7-2 Preference of requests for bandwidth limits

If used bandwidth is		The preference is	Description
Below the minimum		High	Only other High requests delay this request ^a .
Above the minimum	Below the maximum limit.	Medium	High requests are serviced first, then compete with other Medium requests ^a .
	Above the maximum limit, with HARDLIM clear.	Low	Requests are not serviced if any High or Medium requests are available ^a .
	Above the maximum limit, with HARDLIM set.	None	Requests are not serviced.

a. Implementations may occasionally deviate from preference order in servicing requests to meet other goals, such as starvation avoidance.

7.3.7.4 Bandwidth control parameters

The control parameters for bandwidth partitioning schemes are all expressed in a fixed-point fraction of the available bandwidth. See [About the fixed-point fractional format](#).

MPAMCFG_MBW_MAX, the bandwidth control setting register for maximum-bandwidth limit also includes a Hard Limit (HARDLIM) bit that prevents a partition from using more than the maximum fraction of the available bandwidth that is set in that register.

7.3.7.5 Memory-bandwidth allocation accounting window width

For both the minimum- and maximum-bandwidth partitioning schemes, memory-bandwidth regulation occurs over an accounting window. The accounting may be either a moving window or by resetting bandwidth counts at the beginning of each accounting-window period.

The width of the window is discoverable and can be read from `MPAMCFG_MBW_WINWD` for the `PARTID` selected by `MPAMCFG_PART_SEL`.

In implementations that support settable window width per `PARTID`, `MPAMCFG_MBW_WINWD` can be written with a fixed-point format (as described in the register's description) specifying the accounting window width in microseconds.

7.3.7.5.1 Fixed accounting window

In fixed-window accounting, bandwidth is apportioned to requests so that each partition gets bandwidth according to the minimum and maximum for that partition (*Over-allocation of minimum bandwidth*). Request or local priorities (*Priority partitioning*) are used to resolve conflicting requests of the same preference.

When the accounting window's period is reached, a new window begins with no history except for any queue of requests that have not been previously serviced. The new window starts accumulating bandwidth for a partition from zero.

7.3.7.5.2 Moving-window accounting

A moving window tracks partition bandwidth usage by all commands issued in the past window width. There is never a reset of the accounting of bandwidth usage per partition. Instead, bandwidth is added to the accounting when a command is processed and removed from the accounting when that command moves out of the window's history. This continuous accounting is relatively free from boundary effects.

Moving-window accounting requires hardware to track the history of commands within the window, in addition to the bandwidth counters per `PARTID` required by the fixed window.

7.3.7.5.3 Other accounting window schemes

An implementation may use another scheme for maintaining history that is broadly in line with the schemes described here. For example, the current bandwidth might decay at a fixed rate proportional to the bandwidth allocation, but not below a current bandwidth of zero.

7.3.7.6 Over-allocation of minimum bandwidth

The minimum bandwidth allocations of all partitions may sum to more bandwidth than is available. This is not a problem when some partitions are not using their bandwidth allocations, because unused allocations are available for other partitions to use. However, when minimum bandwidth is over-allocated, the minimum bandwidth that is programmed for partitions cannot always be met.

If the programmed minimum bandwidth allocation is to be reliably delivered by the system, software must ensure that minimum bandwidth is not over-allocated.

7.3.7.7 Over-allocation of maximum bandwidth

The maximum bandwidth allocations of all partitions may sum to more bandwidth than is available. This is not a problem when some partitions are not using their maximum bandwidth allocations, because unused allocations are available for other partitions to use. If maximum bandwidth is over-allocated, the maximum bandwidth that is programmed for partitions cannot always be met.

7.3.7.8 Available bandwidth

The bandwidth available downstream from an MSC is not constant, and it affects the operation of minimum and maximum bandwidth partitioning.

Available bandwidth may depend on one or more clock frequencies in many systems (for example, DDR clock). Software may require to reallocate bandwidths when changing clock frequencies that affect available bandwidth. Lowering clock rates without changing allocations may result in over-allocation of bandwidth.

The available bandwidth on a DRAM channel varies with the mix of reads and writes and the bank-hit rate. Bandwidth may also vary with burst size.

7.3.8 Memory-bandwidth proportional-stride partitioning

Proportional-stride bandwidth partitioning control is an instance of proportional resource-allocation generic control, described in *Proportional resource allocation facilities*. The control parameter for bandwidth proportional-stride partitioning is expressed as an unsigned integer.

Regulation according to this scheme permits the partition to consume bandwidth in proportion to its stride, in relation to other requests' strides that are contending for bandwidth. See *Model of stride-based memory bandwidth scheduling* for an example of stride-based proportional bandwidth regulation.

The MPAMF_MBW_IDR.HAS_PROP bit indicates the presence of a memory-bandwidth proportional-stride partitioning control interface in the MSC.

7.3.8.1 Combining memory-bandwidth proportional stride with other memory-bandwidth partitioning

There is no setting of the STRIDEM1 control field that disables the effects of proportional-stride partitioning on a partition's bandwidth usage. To enable proportional-stride partitioning for a PARTID, MPAMCFG_MBW_PROP.EN must be set to 1.

When multiple partitioning controls are active, each affects the partition's bandwidth usage. However, some combinations of controls may not make sense, because the regulation of that pair of controls cannot be made to work in concert.

Memory-bandwidth maximum partitioning must work together with proportional-stride partitioning.

7.3.9 Priority partitioning

Unlike the other memory-system resources in this architecture, priority does not directly affect the allocation of memory-system resources. Instead, it has an effect on conflicts that arise during access to resources. A properly configured system should rarely have substantial performance effects due to prioritization, but priority does play an important role in oversubscribed situations, whether instantaneous or sustained. Therefore, we choose to include priority partitioning here as a tool to aid in isolating memory-system effects between partitions.

A PARTID may be assigned priorities for each component in the memory system that implements a priority partitioning control. This partitioning control allows different parts of the memory system to handle requests with different priorities. For example, requests from a PE to system cache may be set to have a higher transport priority than those from system cache to main memory.

In a system in which the interconnect carries QoS values or priorities, requests arriving at an MSC have an upstream priority as part of the request. In the absence of an internal priority partitioning control, request priority could be used by an MSC to prioritize internal operations. In the absence of a downstream priority partitioning control, the request priority is used as through priority. See *Through priorities*.

Priority partitioning can override the upstream priority with two types of priorities:

- Internal priorities control priorities used in the internal operation of an MSC.
- Downstream priorities control priorities communicated downstream (for example to an interconnect).

“Downstream” refers to the communication direction for requests. “Upstream” refers to the response, and it usually uses the same transport priority as the request that generated it.

7.3.9.1 Internal priorities

Internal priorities are used within an MSC to prioritize internal operations. For example, a memory controller may use an internal priority to choose between waiting requests when bandwidth allocation indicates two or more requests have the same bandwidth preference.

Internal priority partitioning is optional even if downstream priority partitioning is implemented.

7.3.9.2 Downstream priorities

An MSC uses a downstream priority to set transport priorities for downstream requests generated during the servicing of an incoming request from upstream.

Downstream priority partitioning is optional even if internal priority partitioning is implemented.

7.3.9.3 Through priorities

For a system in which the interconnect carries QoS values or priorities, these priorities arrive with incoming requests from upstream. An MSC that does not implement priority partitioning, or that does not implement downstream priority partitioning, must use these upstream priorities on all downstream communication.

If an MSC does not implement priority partitioning, or it does not implement downstream priorities, the downstream priority is always the same as the request (upstream) priority.

The priority of a response through an MSC (from downstream to upstream) is always the same priority as the response received (from downstream). Priority partitioning never alters response priorities received from downstream.

7.4 Vendor or implementation-specific partitioning control interfaces

MPAM provides discoverable vendor extensions to permit partners to invent partitioning controls. These include controls that do not fit the standard interfaces and controls for types of resources not supported through the standard controls defined in this document. Such controls provide product differentiation to address market-segment needs or to provide superior memory-system control.

The `MPAMF_IDR.HAS_IMPL_IDR` bit indicates the presence of `MPAMF_IMPL_IDR` and of implementation-specific or vendor-specific resource partitioning controls.

Vendor, design, or model and version information is present in `MPAMF_IIDR`. `MPAMF_IMPL_IDR` is available for implementations that need to convey additional information about parameters of implementation-specific partitioning controls.

In MPAM v0.1 and from MPAM v1.1:

- If `MPAMF_IMPL_IDR` describes no IMPLEMENTATION DEFINED partitioning controls, `MPAMF_IDR.NO_IMPL_PART` must be 1.
- If `MPAMF_IMPL_IDR` describes no IMPLEMENTATION DEFINED monitors, `MPAMF_IDR.NO_IMPL_MSMON` must be 1.

7.5 Measurements for controlling resource usage

This section is *informative*.

In many cases, resource usage by a partition must be measured so that the resource controller can regulate allocation of the resource to that partition.

In a memory channel, the bytes delivered to requests from a PARTID might be more costly if delivered in response to a series of 1-byte requests rather than cache-line-sized bursts. So, it might be reasonable to count the cost of servicing a 1-byte request to be the same as the cost of servicing a cache-line request rather than as a fraction of a word access cost.

7.6 PARTID narrowing

An implementation may optionally map input PARTID spaces into smaller internal PARTID spaces. This involves mapping the PARTID from a request (reqPARTID) into an internal PARTID (intPARTID). The reqPARTID-to-intPARTID mappings for Secure and Non-secure physical PARTID spaces must be used internally and not for downstream requests.

This mapping is supported by a memory-mapped register, `MPAMCFG_INTPARTID`, and an ID register bit for each of the Secure and Non-secure physical PARTID spaces. The related behavior includes:

- Translate the incoming request's reqPARTID and MPAM_NS into an intPARTID (with the same MPAM_NS) before accessing the control settings and regulation state of the partition.
- Use `MPAMCFG_INTPARTID` to store an association of a reqPARTID in `MPAMCFG_PART_SEL` to the intPARTID stored in `MPAMCFG_INTPARTID`.
- Error code for `MPAMF_ESR` to indicate a bad intPARTID mapping for the reqPARTID.
- A bit in `MPAMCFG_PART_SEL` indicates that the value in that register is an intPARTID. The register can hold either an intPARTID or reqPARTID at any time, but the reqPARTID can only be used for accessing the association by means of `MPAMCFG_INTPARTID`. So, at the time `MPAMCFG_INTPARTID` is read or written, `MPAMCFG_PART_SEL.INTERNAL` must be clear. For access to read or write other control settings registers, the INTERNAL bit must be set.
- With PARTID narrowing implemented, the contents of `MPAMCFG_PART_SEL` are interpreted as an intPARTID for accessing control settings through an `MPAMCFG_*` register other than `MPAMCFG_INTPARTID`. The `MPAMCFG_PART_SEL.INTERNAL` bit must be set to confirm the intPARTID is being used.
- With PARTID narrowing not implemented, the contents of `MPAMCFG_PART_SEL` are interpreted as a reqPARTID. The `MPAMCFG_PART_SEL.INTERNAL` bit must == 0 to confirm that the reqPARTID is being used.

7.7 System reset of MPAM controls in MSCs

This section is *normative*.

After a system reset, the MPAM controls in MSCs must reset the settings for default PARTID so that software can use all of the resource. Since MPAMn_ELx.MPAMEN for the highest implemented ELx is reset to 0 by a system reset, the MPAM fields of all requests issued by a PE use the corresponding default PARTID in the Security state of the PE. Only the resource controls for the default PARTIDs must be reset to full access for the system to behave as if there were no MPAM.

Only the control settings for the default PARTID must be reset. The reset value should be appropriate to allow the default PARTID to access all of the resource. This is needed to allow the system to boot up to a point where MPAM resource controls can be set before non-default PARTIDs are used to make requests.

7.7.1 Suggested reset values for standard control types

Table 7-3 describes the suggested reset values for PARTID == 0 control setting for both MPAM_NS == 0 and MPAM_NS == 1.

Table 7-3 Suggested reset values for PARTID0 controls

Control type	Reset value
MPAMCFG_CPBM<n>	All ones for all implemented <i>n</i>
MPAMCFG_CMAX.CMAX	0xFFFF
MPAMCFG_CMAX.SOFTLIM	0xb0
MPAMCFG_MBW_PBM<n>	All ones for all implemented <i>n</i>
MPAMCFG_MBW_MAX.MAX	0xFFFF
MPAMCFG_MBW_MAX.HARDLIM	0b0
MPAMCFG_MBW_MIN.MIN	0xFFFF
MPAMCFG_MBW_PROP	EN=0
MPAMCFG_CASSOC.CASSOC	0xFFFF
MPAMCFG_EN_FLAGS.EN0	0b1

In addition, for PARTID narrowing, Arm recommends that reqPARTID == 0 map to intPARTID == 0 and that the reset values be applied to the settings of intPARTID == 0 in both values of MPAM_NS.

7.8 About the fixed-point fractional format

This section is *normative*.

Fractional control parameters use a 16-bit fixed-point format. The format permits implementations to have fewer than 16 bits by truncating least significant bits from the fraction and implementing these bits as RAZ/WI.

Software can be expected to calculate a 16-bit fractional part to store into the memory-mapped register without the need to understand the implemented width of the field. If the field width is less than 16 bits, the least significant bits are silently IGNORED by the implementation. This results in an uncertainty of the intended value.

If software stores an intended fractional value into a field with an implemented width of w , the implementation's truncated field sees a value of v . The value v is at the bottom of the range of v to $v + 2^{-w} - 2^{-17}$ and the intended fractional value lies somewhere within that range, inclusive of the end points.

Depending on the use of the fractional value, the best choice of value within the range could be the center of the range, the smallest end of the range, or the greatest end of the range. For examples, a cache maximum-capacity fraction might best be interpreted as the highest end of the range, and a cache minimum-capacity fraction might best be interpreted as the lowest end of the range.

Table 7-4 shows the fraction widths and hex representation used for three formats. The values in the table are suitable for a maximum limit because the Max value for every entry is never greater than the target value.

Table 7-4 Fraction Widths and Hex Representation

Percentage	16 bits			12 bits			8 bits		
	Hex	Min	Max	Hex	Min	Max	Hex	Min	Max
1.00%	028E	0.9979%	0.9995%	027	0.9521%	0.9766%	01	0.3906%	0.7813%
12.50%	1FFF	12.4985%	12.5000%	1FF	12.4756%	12.5000%	1F	12.1094%	12.5000%
16.67%	2AAB	16.6672%	16.6687%	2A9	16.6260%	16.6504%	29	16.0156%	16.4063%
25%	3FFF	24.9985%	25.0000%	3FF	24.9756%	25.0000%	3F	24.6094%	25.0000%
33.33%	5552	33.3282%	33.3298%	554	33.3008%	33.3252%	54	32.8125%	33.2031%
35%	5998	34.9976%	34.9991%	598	34.9609%	34.9854%	58	34.3750%	34.7656%
37.25%	5F5B	37.2482%	37.2498%	5F4	37.2070%	37.2314%	5E	36.7188%	37.1094%
42.50%	6CCB	42.4973%	42.4988%	6CB	42.4561%	42.4805%	6B	41.7969%	42.1875%
45%	7332	44.9982%	44.9997%	732	44.9707%	44.9951%	72	44.5313%	44.9219%
50%	7FFF	49.9985%	50.0000%	7FF	49.9756%	50.0000%	7F	49.6094%	50.0000%
52%	851D	51.9974%	51.9989%	850	51.9531%	51.9775%	84	51.5625%	51.9531%
55%	8CCB	54.9973%	54.9988%	8CB	54.9561%	54.9805%	8B	54.2969%	54.6875%
58%	9479	57.9971%	57.9987%	946	57.9590%	57.9834%	93	57.4219%	57.8125%
62.75%	A0A2	62.7472%	62.7487%	A09	62.7197%	62.7441%	9F	62.1094%	62.5000%
66.67%	AAA9	66.6641%	66.6656%	AA9	66.6260%	66.6504%	A9	66.0156%	66.4063%
75%	BFFF	74.9985%	75.0000%	BFF	74.9756%	75.0000%	BF	74.6094%	75.0000%
82.50%	D332	82.4982%	82.4997%	D32	82.4707%	82.4951%	D2	82.0313%	82.4219%
88%	E146	87.9974%	87.9990%	E13	87.9639%	87.9883%	E0	87.5000%	87.8906%
95%	F332	94.9982%	94.9997%	F32	94.9707%	94.9951%	F2	94.5313%	94.9219%

Table 7-4 Fraction Widths and Hex Representation (continued)

Percentage	16 bits			12 bits			8 bits		
	Hex	Min	Max	Hex	Min	Max	Hex	Min	Max
100%	FFFF	99.9985%	100.0000%	FFF	99.9756%	100.0000%	FF	99.6094%	100.0000%
2^n	65536			4096			256		
ndigits	4			3			2		
shift	0			0			0		

Chapter 8

Resource Monitors

This chapter contains the following sections:

- *Introduction.*
- *MPAM resource monitors.*
- *Common features.*
- *Monitor configuration.*
- *Monitor behavior on overflow.*

8.1 Introduction

Software environments may be labeled as belonging to a Performance Monitoring Group (PMG) within a partition. The PARTID and PMG can be used to filter some performance events so that the performance of a particular PARTID and PMG can be monitored.

8.2 MPAM resource monitors

MPAM resource monitors provide software with measurements of the resource-type usage that can be partitioned by MPAM. There are two types of MPAM resource monitors:

- *Memory-bandwidth usage monitors*
- *Cache-storage usage monitors*

Each type of monitor measures the usage by memory-system transactions of a PARTID and PMG. An MSC may implement any number of performance monitor instances, up to 2^{16} of each type. The PARTID for filtering resource monitors is always a request PARTID, even when PARTID narrowing is implemented.

To access a monitor instance, the instance number is stored into the `MSMON_CFG_MON_SEL.MON_SEL` field. All of the monitor access registers for a type of monitor then access that instance of that type. See *Monitor configuration*.

If the implementation supports the RIS MPAM feature, the MSC may have two or more partitionable resources differentiated by the value of Resource Instance Selector (RIS). See *Resource instance selection*.

MPAM resource monitors are associated with a particular MPAM partitionable resource, but memory bandwidth monitors may be placed at the top level of an MSC. Monitors at the top level of an MSC are accessed with the RIS value of 0.

The monitor instance accessed by the MSC's `MSMON_*` registers is controlled by setting `MSMON_CFG_MON_SEL.MON_SEL` to the instance number to access. If the implementation supports the RIS MPAM feature, the MSC may have two or more partitionable resources differentiated by the value of `MSMON_CFG_MON_SEL.RIS`.

8.2.1 Memory-bandwidth usage monitors

A memory-bandwidth usage monitor counts payload bytes meeting the filter criteria that pass the monitoring point in the downstream direction for writes or the upstream direction for reads. Each monitor has the following set of memory-mapped configuration registers and functional features:

- A control register `MSMON_CFG_MBWU_CTL` that configures behavior of the monitor instance.
- A filter register `MSMON_CFG_MBWU_FLT` that specifies the transfers to be counted. This register has fields for reads, writes, PARTID, PMG, and other criteria.
- A monitor register `MSMON_MBWU` that contains an optionally scaled count of bytes transferred downstream from this MSC that match the conditions of the filter register. This monitor register may be reset after each capture event. If scaling is enabled, the value read from `MSMON_MBWU` must be shifted left by `MPAMF_MBWUMON_IDR.SCALE` bit positions to scale the value to the number of bytes.
- In MPAM v0.1 and from MPAM v1.1, an optional long monitor register, `MSMON_MBWU_L`, that contains a count of 44 bits or 63 bits. A NRDY bit is also present in this register, see *Not-ready Bit*.
- An optional capture register `MSMON_MBWU_CAPTURE` that is loaded from the monitor register each time the selected capture event occurs. When a capture event occurs, the monitor register is copied to the capture register and the monitor register is optionally reset to zero.
- In MPAM v0.1 and from MPAM v1.1, if `MPAMF_MBWUMON_IDR.{HAS_LONG, HAS_CAPTURE}` are 1, the `MSMON_MBWU_L_CAPTURE` register must be implemented.
- A Not-Ready (NRDY) bit (*Not-ready Bit*) in the memory-bandwidth usage register `MSMON_MBWU` is set when the filter register or the control register is written. The NRDY bit is reset to 0 after a capture event. The NRDY bit is copied to the capture register along with the rest of the monitor register's content. This copy is made before the NRDY bit is reset. If the value of the NRDY bit in the capture register is 1, the captured resource usage should be viewed as representing an incomplete sampling interval. Therefore, the count should be assumed to be incorrect.

A capture event is needed if the optional capture register is implemented. The capture event causes the transfer of each monitor's count register to its capture register and may optionally reset the count register.

If the count register is reset by a capture event, this allows reading the bytes transferred that meet the criteria set in the filter and control registers:

- During the interval between the last two capture events from `MSMON_MBWU_CAPTURE`.
- Since the last capture event from `MSMON_MBWU`.

Bandwidth usage can be computed in software from the count of bytes transferred as read from `MSMON_MBWU` or `MSMON_MBWU_CAPTURE` and the interval over which the count was collected.

There can be several sources of the capture event. The capture event source to use is specified in `MSMON_CFG_MBWU_CTL.CAPT_EVNT` (*Memory-mapped monitoring configuration registers*). It can be advantageous to use a single event to capture monitors in several MSCs simultaneously. A periodic capture event for multiple MSCs could be generated at the system level, perhaps using a generic timer, and distributed to the MSCs.

The source of an external capture event is selected in `MSMON_CFG_MBWU_CTL.CAPT_EVNT`. A local capture event generator is present if `MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT == 1`, and this generator generates events when certain values are written into `MSMON_CAPT_EVNT`.

If `MSMON_CFG_MBWU_CTL.OFLOW_LNKG` is implemented, a monitor instance with an `OFLOW_LNKG` value that is not zero can signal the capture event number in that field when the monitor instance overflows. See *Control of signaling to other monitor instances*.

8.2.1.1 Scaled MBWU count value

If `MSMON_CFG_MBWU_CTL.SCLEN == 0`, the count is not scaled. If `MSMON_CFG_MBWU_CTL.SCLEN == 1`, the count in `MSMON_MBWU` is a scaled count.

The scaled count in `MSMON_MBWU` is the true count of bytes transferred, rounded to 2^{SCALE} and then shifted right by `SCALE` bit positions. The shift count, `SCALE`, is `MPAMF_MBWUMON_IDR.SCALE`.

`SCALE` is an implementation constant chosen for a monitoring point such that periodic sampling and reset of `MSMON_MBWU_CAPTURE` can count the highest traffic rates possible at the monitoring point without overflowing the `VALUE` field at a maximum sampling rate. The sampling rate is limited by the target use.

For example, if the maximum traffic that could pass the monitoring point is 300 Gbps and the system environment supports capturing the counter 30 times per second, the counter must be scaled to no more than $2^{31} - 1$ counts per thirtieth of a second. This requires scaling the counter by a factor of at least 5, so the `SCALE` must be at least 3.

If the traffic to memory might be distributed across several MSCs (for example, across several memory channel controllers), a comprehensive measurement of bandwidth might require reading multiple memory-bandwidth usage monitors on those MSCs and summing the results. Capturing those monitors with the same system-level capture event allows correlated monitor values.

8.2.1.2 Long MBWU counter and capture

In MPAM v0.1 and from MPAM v1.0, there is optional support for 44-bit or 63-bit MBWU counters.

`MSMON_MBWU_L` is optional and only present when `MPAMF_MBWUMON_IDR.HAS_LONG` is 1. This indicates that this monitor type supports long counters.

If `MPAMF_MBWUMON_IDR.HAS_LONG` and `MPAMF_MBWUMON_IDR.HAS_CAPTURE` are 1, the `MSMON_MBWU_L_CAPTURE` register must also be implemented.

The `VALUE` field of the long registers is never scaled.

The `VALUE` field of `MSMON_MBWU_L` and `MSMON_MBWU_L_CAPTURE` can be implemented either as a 63-bit `VALUE` field or a 44-bit `VALUE` field. The 44-bit `VALUE` field is indicated when `MPAMF_MBWUMON_IDR.LWD` is 0 and has bits[62:44] of each register as `RES0`. When `MPAMF_MBWUMON_IDR.LWD` is 1, the `VALUE` field of each register is 63 bits.

An overflow occurs in the long counter when the count in the VALUE field exceeds the maximum representable value. This depends on the length of the VALUE field. In `MSMON_MBWU`, the VALUE field is always 31 bits. If `MSMON_MBWU_L` is implemented, the length of the VALUE field is either 63 or 44 bits as set by `MPAMF_MBWUMON_IDR.LWD`.

When any instance of the `MSMON_MBWU_L` counter overflows, the `MSMON_CFG_MBWU_CTL.OFLOW_STATUS_L` bit is 1. If `MSMON_CFG_MBWU_CTL.OFLOW_INTR_L` is set, this overflow produces an MPAM Overflow interrupt. See *MPAM overflow interrupt* and *Control of monitor behavior on overflow*.

When an implementation has both the long counter and the short 31-bit counter, the short counter might overflow when the long counter has not overflowed and produce an MPAM Overflow interrupt. This can be prevented by setting `MSMON_CFG_MBWU_CTL.OFLOW_INTR` to 0, which disables the overflow interrupt for overflow of the short counter.

The `MSMON_CFG_MBWU_CTL.OFLOW_FRZ` field is not duplicated, and affects the behaviors of both short and long counters on overflow.

8.2.2 Cache-storage usage monitors

A cache-storage usage monitor is filtered by a PARTID and PMG. Each monitor has the following memory-mapped configuration registers:

- A filter register `MSMON_CFG_CSU_FLT` that sets the PARTID and PMG to be monitored.
- A cache-storage usage register `MSMON_CSU` that reports the amount of storage currently present within the cache allocated by the PARTID and PMG. It is an implementation choice whether `MSMON_CSU` is implemented as RO or RW.
- A Not-Ready bit in the cache-storage usage register `MSMON_CSU` that indicates that the value is not accurate. An implementation may set this NRDY bit if the value in the cache-storage usage register is not currently accurate, possibly because it is still being computed. For more information on the Not-Ready bit, see *Not-ready Bit*.
- An optional capture register `MSMON_CSU_CAPTURE` that is loaded from the cache-storage usage register each time the capture event occurs.

A capture event is needed if the optional capture register is implemented. The capture event causes the transfer of each monitor's cache-storage usage register to its optional capture register.

The source of an external capture event is not specified here. It can be advantageous to use a single event to capture monitors in several MSCs simultaneously. A periodic capture event for multiple MSCs could be generated at the system level, perhaps using a generic timer, and distributed to the several MSCs.

The source of an external capture event is selected in `MSMON_CFG_CSU_CTL.CAPT_EVNT`. A local capture event generator is present if `MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT == 1`, and this generator generates events when certain values are written into `MSMON_CAPT_EVNT`.

If a monitor needs time to become accurate, the NRDY bit signals that the value is not yet accurate. Some methods of building cache-storage usage monitors might involve (1) a phase in which the monitor collects enough information to begin accurately tracking usage, or (2) a phase in which the measurement is kept accurate by tracking resource usage events. For example such a monitor might take tens of microseconds to complete the first phase before the value accurately tracks the actual resource usage. In this case, the NRDY bit would be kept at 1 until the monitor value becomes accurate.

The NRDY bit is included because some implementations may have timing restrictions between setting the filter register and reading the cache-storage usage register that may span thousands of PE cycles. Reading the monitor too soon is permitted to affect the accuracy of the readout, and it is indicated when the NRDY bit of the cache-storage usage register is 1.

The cache storage usage monitor architecture supports overflow behavior in CSU monitors. However, Arm recommends that CSU monitors be designed so that overflow is not possible.

8.3 Common features

All MPAM performance monitors have these features:

- Monitor register.
- Not-ready bit.
- Capture register.
- Overflow bit.
- Enable bit.

These features are described below.

8.3.1 Monitor register

Every monitor instance has a monitor register that contains the VALUE field and the NRDY field. The VALUE field contains the current count or measurement value of the monitor.

When the monitor is enabled, the VALUE field can change at any time. If this monitor counts events, such as the memory bandwidth usage monitor counts bytes passing a monitoring point, the count increases as monitored events occur that match the filter criteria. Software can reset the VALUE field to have this monitor count register produce a count based entirely on the updated field.

A monitor that measures resource usage, such as a cache storage usage monitor, measures the bytes in cache lines of a cache. This gives the automatically measured usage in its VALUE field. The measurement can move up and down as the resource usage, thereby matching the filter criteria changes.

Data written to the monitor register by software would be overwritten by the automatically measured value.

The implementer of a measurement monitor can choose to implement the monitor register as read-only. For example, see [MPAMF_CSUMON_IDR.CSU_RO](#).

The VALUE field of a measurement monitor register could have an initial period where it is temporarily inaccurate while converging on the measurement. During this period, the NRDY bit is 1 to indicate an inaccurate measurement. See *Not-ready Bit*.

It is not mandatory to disable a monitor to reprogram it. However, if both [MSMON_MBWU](#) and [MSMON_MBWU_L](#) are present in an MSC, the VALUE fields of the two registers cannot be reset simultaneously. In this case the recommended software flow is:

1. Set [MSMON_CFG_MON_SEL](#) with the desired monitor instance selector and RIS.
2. Set the [MSMON_CFG_MBWU_CTL.EN](#) bit to 0.
3. Reconfigure both monitors by setting their VALUE fields to 0.
4. Change the filter settings of these registers in [MSMON_CFG_MBWU_FLT](#).
5. Set [MSMON_CFG_MBWU_CTL](#) to the new configuration settings and set the EN field to 1.

The events counted or the resources measured by a monitor instance depend upon the configuration of the monitor instance's configuration and filter registers. If either of these are changed, the VALUE field continues counting or measuring with the new criteria without automatically resetting VALUE to 0.

A counting monitor register can be set by software to any value by storing the new value to the register. Counting monitors continue counting the monitored event from this new value.

A write to a monitor register sets the NRDY and VALUE fields to the data written. The monitor register can be written at any time.

When [MSMON_CFG_<type>_CTL.EN](#) is 1, the monitor is Enabled to count events or measure resource usage matching the configuration in [MSMON_CFG_<type>_CTL](#) and [MSMON_CFG_<type>_FLT](#). An Enabled monitor automatically updates as matching events occur or matching resource usage changes.

When `MSMON_CFG_<type>_CTL.EN` is 0, the monitor is Disabled. A Disabled monitor must not change automatically, but may be written by software.

A monitor overflows when the `VALUE` field exceeds its largest representable value. After an overflow, the `VALUE` field wraps around by dropping the high-order bit. The resulting truncated value may be zero or greater, depending on the increment for a counter or the value for a measurement.

When an overflow occurs, the overflow behaviors selected in `MSMON_CFG_<type>_CTL` are performed automatically. These independently-selectable behaviors are:

- Freezing the `VALUE` field or continuing to count or measure.
- Signaling an MPAM overflow interrupt from the MSC.
- Capturing the overflowed `VALUE` field in the monitor instance's capture register and optionally resetting the monitor's `VALUE` and `NRDY` fields to zeros.
- Signaling a capture event to other monitor instances that are programmed to be sensitive to that capture event or different types of and for the same or different RIS numbers.

When an overflow occurs, the overflow status bit in `MSMON_CFG_<type>_CTL` must be set to 1.

8.3.2 Not-ready Bit

The Not-ready (`NRDY`) bit, in the `MSMON_MBWU` and `MSMON_CSU` registers, when set, indicates that the monitor does not have an accurate count or measurement yet, because the monitor's settings have been recently changed. If the monitor requires some time to establish a new count or measurement after its settings are changed, the `NRDY` bit must be set automatically when the settings are changed and reset when the count or measurement is accurately represented in the monitor.

In the absence of another change in settings, the `NRDY` bit must clear automatically within a maximum length of time. The maximum time that `NRDY` may be 1 is an implementation parameter that is discoverable in the firmware data value of `MAX_NRDY_USEC` for the MSC's monitor type. For example, a measurement slow to respond to changes to what it measures could take up to `MAX_NRDY_USEC` to converge to a new fresh measurement.

Each instance of each type of monitor keeps its `NRDY` bit separately. For example, if `MBWU` monitor instance 3 is collecting memory bytes transferred for one partition and `MBWU` monitor instance 6 is later configured to collect for another partition, the configuration of `MBWU` monitor instance 6 must not disturb the on-going collection in `MBWU` monitor instance 3.

The `NRDY` bit of a monitor or capture register can be written to either state. On a monitor that measures resource and requires time to reach an accurate value, the `NRDY` bit must automatically reset when the measurement has become accurate. On a counting monitor, the `NRDY` bit remains set until it is reset by software writing it as 0 in the monitor register, or automatically after the monitor is captured in the capture register by a capture event.

If a monitor supports the automatic behaviors of `NRDY`, it must clear the `NRDY` when its measurements are accurate. The monitor must also clear the `NRDY` if it is configured for capture, after a capture event causes transfer to the capture event register.

If a monitor does not support automatic behavior of `NRDY`, software can use this bit for any purpose.

8.3.3 Capture event and capture register

Fields in `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL` control the behavior of a monitor instance that receives the capture event. Both registers must have the same fields for all monitor instances of a resource instance and monitor type.

A capture event causes every monitor that is configured to be sensitive to that event to be copied into that monitor's capture register.

Capture events may be local to the MSC or external to the MSC and may be software-initiated single events or a periodically repeating series of events. External capture events are system-defined. A generic counter can be used as the source of such an event, but this is not required. An external capture event could be distributed to all MSCs so that system-wide captures occur of all monitors sensitive to the external event. This permits using the various measurements for sums and differences because they measure the same period and (mostly) related resource usage.

A capture register for a monitor is loaded with the monitor's count or measurement and its NRDY bit when a capture event that is selected in the monitor's control register occurs. A capture event completes almost instantaneously, so no handshake is used for completion. However, the NRDY bit indicates whether a capture is not an accurate reading.

If the event is periodic, software can read the capture registers at any time to get the results captured when the most recent capture event occurred.

If it makes sense for the particular monitored value, the count or measurement can optionally be reset by the event. In this case, the value in the capture register represents a count over the capture-event period or a measurement over that period.

There are eight capture event codes:

Table 8-1 Capture events code and function

Capture event code	Function
0	No capture event
1 - 6	Available for use
7	Local capture event

A monitor instance does not respond to any capture event when the CAPT_EVNT field of the control register is set to 0, no capture event. This field can be set to an event number of 1 through 7. The instance is sensitive to a capture event matching the event number that CAPT_EVNT is set to.

Capture events are available to all monitors in the MSC. All monitor instances in all resource instances and of all monitor types can each be sensitive to any of the 7 capture event codes. Every monitor instance monitoring a PARTID space must be able to be sensitive to the same capture event.

8.3.3.1 Local capture-event generator

If `MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT == 1`, the `MSMON_CAPT_EVNT` register exists and generates capture events that are local to an MSC when it is written with a value that contains a 1 in the NOW bit position.

There are separate `MSMON_CAPT_EVNT` registers for Secure and Non-secure address spaces. The Non-secure version generates a local capture event to all Non-secure monitors within the MSC that have been configured to use `MSMON_CFG_<type>_FLT.CAPT_EVNT == 7` (Table 8-2). The Secure version of `MSMON_CAPT_EVNT` generates a local capture event to all Secure monitors within the MSC that have been configured to use `CAPT_EVNT == 7` when `MSMON_CAPT_EVNT` is written with `ALL == 0` and `NOW == 1`. When the ALL and NOW bits both == 1 in a write to Secure `MSMON_CAPT_EVNT`, the write generates a local capture event to all Secure and Non-secure monitors within the MSC that have been configured to use `CAPT_EVNT == 7`.

If `MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT == 0`, local capture events are not generated and any monitors that have their control register set to `CAPT_EVNT == 7` do not receive any capture events.

8.3.3.2 Reset on capture

Monitors that keep a count of events, or that accumulate counts such as bytes transferred, may be optionally reset after a capture event transfers the count to the monitor's capture register. This behavior on capture is controlled by the `MSMON_CFG_*_CTL.CAPT_RESET` bit. If `CAPT_RESET == 1`, the monitor count is reset to 0 immediately after the value is captured into the `MSMON*_CAPTURE` register.

Monitors that report a current resource value, such as cache-storage usage, that cannot reasonably be reset, do not need to support reset on capture behavior. Arm recommends that these monitors have the `CAPT_RESET` bit as RAZ/WI.

8.3.4 Overflow status bit

The `MSMON_CFG_<type>_CTL.OFLOW_STATUS` bit is set to 1 when the monitor counter overflows. This bit must be reset by writing 0 to the `OFLOW_STATUS` field.

The `MSMON_CFG_<type>_CTL` register contains fields to control MPAM behavior on an overflow. The `OFLOW_FRZ` bit, when set, freezes the counter after the count that caused it to overflow. When reset to 0, the counter continues to count after an overflow.

If the overflow changes the `OFLOW_STATUS` flag from 0 to 1 and the `OFLOW_INTR` bit is set, an MPAM overflow interrupt will be signaled if implemented. See also *MPAM overflow interrupt*.

8.3.5 Enable bit

The `MSMON_CFG_<type>_CTL.EN` bit is set to 1 when the monitor is enabled to collect information according to its configuration. When the `EN` bit is 0, the monitor is disabled and must not count events or measure resources. The monitor configuration registers can be written and read, regardless of the value of the monitor `EN` bit field.

8.4 Monitor configuration

For each type of resource monitor, the number of monitor instances that are available is described in the corresponding MPAMF_<type>MON_IDR.NUM_MON field.

The MSMON_CFG_MON_SEL.MON_SEL field selects the monitor instance to configure. The MON_SEL monitor instance of monitor type, type, is accessed when an MSMON_CFG_<type> register is accessed.

All monitor types have two 32-bit configuration registers:

- MSMON_CFG_<type>_FLT (Table 8-2) has fields to select the PARTID and PMG to monitor.
- MSMON_CFG_<type>_CTL (Table 8-2) has controls for counting a subset of events, controlling overflow, and capture behavior.

Some monitor types may not require all fields, and fields not required must be RAZ/WI or RAO/WI.

8.5 Monitor behavior on overflow

When an MPAM monitor instance overflows the OFLOW_STATUS flag in its configuration register is set to 1. Each monitor instance can be configured for a number of optional behaviors in its configuration register:

- Broadcast a configured capture event number to all other monitors in the MSC. See *Control of signaling to other monitor instances*.
- Freeze a monitor instance. See *Control of monitor behavior on overflow*.
- If the capture register is implemented the monitor instance can be captured to its capture register.
- Signal an overflow interrupt. See *Control of monitor behavior on overflow* and *MPAM overflow interrupt*.

A monitor instance can be configured to respond to a capture event as an overflow. A capture event configured to be received by other monitor instances as an overflow can be caused by the overflow of a monitor that is configured to broadcast the capture event when the monitor overflows.

A group of monitor instances can be configured so that if any of them overflows, the overflow signals the others on a particular capture event number and all of the others respond by also performing overflow behaviors. See *Control over behavior of a monitor instance on a capture event*.

Capture events are local to the MSC and are not broadcast to other MSCs.

There are seven capture event codes. See *Capture events code and function*.

8.5.1 Control of monitor behavior on overflow

The behavior of a monitor instance on overflow is governed by fields in the `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL` registers:

Table 8-2 Control on overflow

Field	Bit	Control of behavior
OFLOW_LNKG	[10:8]	A capture event may be signaled to other monitor instances within the MSC with the capture event number in this field if it is not zero.
OFLOW_CAPT	[23]	The monitor instance VALUE field is captured into its capture register after the overflow occurs.
OFLOW_FRZ	[24]	The monitor instance VALUE field does not change after it overflows.
OFLOW_INTR	[25]	An interrupt is signaled after the monitor instance OFLOW_STATUS changes from 0 to 1.

The OFLOW_STATUS field of the `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL` registers for a monitor instance is set when the monitor instance VALUE overflows.

The monitor does not change the VALUE field to count events or update the resource measurement when it is frozen. The monitor instance resumes counting or measuring when OFLOW_STATUS is reset, either by writing `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL` registers with the OFLOW_STATUS field as 0 or by writing the monitor count register.

Writing the monitor count register resets OFLOW_STATUS as software often writes the monitor count register with a new starting value, for example 0.

8.5.2 Control of signaling to other monitor instances

The field OFLOW_LNKG in `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL` controls signaling other monitor instances when the monitor instance overflows.

If the monitor implements overflow linkage as indicated in `MPAMF_<type>_MON_IDR.HAS_OFLOW_LNKG` as 1, the monitor instances can be configured to signal a capture event when the count or measurement in a monitor register overflows.

To signal the overflow linkage capture event on the overflow of the VALUE field of an instance of `MSMON_<type>`, all of the following conditions must be met in the `MSMON_CFG_<type>_CTL` register:

- The `OFLOW_LNKG` field must be set to a non-zero value corresponding to a capture event implemented in the MSC.
- The `OFLOW_INTR` bit must be set to 1.

To signal the overflow linkage capture event on the overflow of the VALUE field an instance of `MSMON_<type>_L`, all the following conditions must be met in the `MSMON_CFG_<type>_CTL` register:

- The `OFLOW_LNKG` field must be set to a value other than zero, corresponding to a capture event implemented in the MSC.
- The `OFLOW_INTR_L` bit must be set to 1.

The capture event indicated in the `MSMON_CFG_<type>_CTL.OFLOW_LNKG` field is to all the monitor instances in the MSC that monitor the same PARTID space as the instance signaling the overflow.

8.5.3 Control over behavior of a monitor instance on a capture event

An external capture event is distributed to monitors for all PARTID spaces of the of the MSC.

Local capture events are distributed to:

- Monitor instances of the Non-secure PARTID space if originated from the Non-secure address space.
- Monitor instances of the Secure address space if originated from the Secure address space.
- Monitor instances of the Root PARTID address space if originated from the Root address space.
- Monitor instances of the Real PARTID address space if originated from the Realm address space.
- Monitor instances for both the Secure and Non-secure PARTID spaces if originated from the Secure address space that has the `MSMON_CAPT_EVNT.ALL` bit set to 1.
- Monitor instances for the same PARTID space as the overflowing monitor instance if the capture events are raised by overflow of a monitor instance with `OFLOW_LNKG` set to 1 through 6.

The behavior of a monitor instance that receives a capture event to which it is sensitive is controlled by settings in its control register. See `MSMON_CFG_CSU_CTL` and `MSMON_CFG_MBWU_CTL`.

8.5.3.1 Configuring the handling capture events as linked overflows

A capture event can be handled by a monitor instance as an overflow linkage and processed much as if this monitor instance had overflowed.

To configure a capture event to be handled as an overflow linkage, the `MSMON_CFG_<type>_CTL.CEVNT_OFLW` bit is set to 1 and `CAPT_EVNT` is set to the capture event number to be used for overflow linkage. The following controls in `MSMON_CFG_<type>_CTL` control the processing of a capture event when `CEVNT_OFLW` is 1:

- `OFLOW_FRZ` [24]: The monitor instance's VALUE is frozen and does not change until the monitor's VALUE has been written. If the monitor implements both normal and long versions of the count, both are frozen and each must be written to unfreeze its VALUE.
- `OFLOW_CAPT` [23]: The monitor instance's VALUE and NRDY fields are copied to the capture register. This field is present when `MPAMF_<type>MON_IDR` fields `HAS_CAPTURE` and `HAS_OFLOW_CAPT` are both 1.
- `OFLOW_CAPT_L` [13]: The monitor instance's long VALUE and NRDY fields are copied to the long capture register. This field is only in `MSMON_CFG_MBWU_CTL`. This field is present when `MPAMF_<type>MON_IDR` fields `HAS_CAPTURE`, `HAS_LONG` and `HAS_OFLOW_CAPT` are all 1.
- `CAPT_RESET` [27]: After the monitor instance is copied to its capture register, the monitor is reset to zero.

8.5.4 Monitors with and without capture

If capture is implemented for the monitor type, the following behavior is followed if only some of the control fields are implemented:

- Overflow linkage for signaling other monitor instances: A monitor that does not implement capture events can signal instances of other monitors that do if it implements OFLOW_LNKG and OFLOW_INTR in MSMON_CFG_<type>_CTL. If the long monitor register is also implemented, MSMON_CFG_MBWU_CTL.OFLOW_INTR_L must also be implemented to control signaling of an overflow linkage event when the long monitor register's VALUE field overflows.

If capture is not implemented for the monitor type, the following behavior is followed if only some of the control fields are implemented:

- Freezing a monitor instance on a capture event: If capture is not implemented, a monitor instance that has not overflowed can be frozen by a capture event if it implements CAPT_EVNT and OFLOW_FRZ. To achieve this behavior, CAPT_EVNT must be set to a capture event 1 through 7. OFLOW_FRZ must be 1.

Chapter 9

Memory-mapped Registers

This chapter contains the following sections:

- *Overview of MMRs.*
- *Summary of memory-mapped registers.*
- *Memory-mapped ID register description.*
- *Memory-mapped partitioning configuration registers.*
- *Memory-mapped monitoring configuration registers.*
- *Memory-mapped control and status registers.*

9.1 Overview of MMRs

The MPAM behavior of an MSC is discovered and configured through memory-mapped registers (MMRs) in the MSC.

All MPAM MMRs are located on one of the MPAM feature pages for the MSC (*MPAM feature page*). An MSC's MPAM feature page is located from information about the device, possibly provided via firmware data such as device tree or ACPI (*Appendix B MSC Firmware Data*).

An MPAM feature page exists in the Non-secure address space and another exists in the Secure address space. If FEAT_RME is supported there is also a Root MPAM feature page in the Root address space and a Realm MPAM feature page in the Realm address space.

The addresses of the MPAM feature pages of an MSC do not need to have the same base address. Arm recommends that the numerical base addresses of the MPAM feature pages in different address spaces be sufficiently different that the numerical address ranges do not overlap.

MPAM MSC MMRs must support 32-bit access as a single access. There is no requirement that accesses of wider than 32 bits complete atomically.

There are MMRs for identifying MPAM parameters and options, the ID registers. These IDRs have the MPAMF prefix.

Other registers configure MPAM resource controls. These registers have the MPAMCFG prefix.

The resource monitor configuration and readout registers have the MSMON prefix.

There is a register to report the status of MPAM programming errors encountered in the MSC and a register to control MPAM interrupts.

9.1.1 Determining presence and location of MMRs

The `MPAMF_IDR` register is located at offset `0x0000` of the MPAM feature page. It indicates which MPAM resource controls are present in the MSC and the maximum PARTID and PMG supported in requests to the MSC. Other MPAMF ID registers are present if the corresponding `MPAMF_IDR` register bit is set and those registers identify the implemented values of architecturally-defined parameters associated with the particular class of MPAM resource control.

The `MPAMF_IDR` also indicates whether the MSC has MPAM monitors. If so, `MPAMF_MSMON_IDR` indicates which monitor types are supported by the MSC. Other monitor MPAMF ID registers are present if the corresponding bit in `MPAMF_MSMON_IDR` is set and those registers identify the implemented values of architecturally-defined parameters associated with the particular type of MPAM monitor.

The address of each MPAM MMR present in an MSC is located within the MPAM feature page for that component at a register-specific offset into that page. The offsets are given in tables in *Summary of memory-mapped registers* and *MPAM feature page*.

9.1.2 Configuring resource controls for a partition

To configure the MPAM resource controls supported by an MSC for a PARTID:

1. Gain exclusive access to the MSC's partitioning configuration registers (for example, take a lock for the memory-mapped partitioning configuration registers, *Memory-mapped partitioning configuration registers*).
2. Write the PARTID to the component's `MPAMCFG_PART_SEL`.
3. Write to the `MPAMCFG_*` registers for the resource controls of the component.
4. Repeat step 3 to configure additional controls associated with the PARTID selected in step 2.
5. Repeat steps 2 through 4 to configure controls for additional PARTIDs.
6. Release exclusive access to the MSC's partitioning control configuration registers (for example, release the lock taken in step 1).

Repeat this procedure for each MSC.

The configuration registers are all the read-write registers that begin with MPAMCFG_*. That is all of the registers in *Memory-mapped partitioning configuration registers*. Before writing any of these registers, software must take a lock to prevent other software from accessing these registers until the lock is released. This is in part because the writing involves first putting a PARTID into the MPAMCFG_PART_SEL register and then writing a configuration value into one or more of the MPAM resource control's configuration registers (also MPAMCFG_* registers).

Software must also take a lock to read any MPAMCFG_* register, other than MPAMCFG_PART_SEL, because reading also involves first putting a PARTID into MPAMCFG_PART_SEL register and then reading a configuration value from one or more of the MPAMCFG_* registers.

There are two copies of MPAMCFG_PART_SEL, one for resource controls for the Secure PARTID space that are accessed from the Secure address space, and the other for resource controls for the Non-secure PARTID space that are accessed from the Non-secure address space. Because there are two copies, there can be separate locks for Secure MPAMCFG_PART_SEL and for Non-secure MPAMCFG_PART_SEL.

9.1.3 Configuring memory-system monitors

To configure the memory-system monitors supported by an MSC for a PARTID and PMG:

1. Gain exclusive access to the MSC's monitor configuration registers (for example, take a lock for the memory-mapped monitoring configuration registers, *Memory-mapped monitoring configuration registers*).
2. Write to the component's MSMON_CFG_MON_SEL to select one of the monitor instances available in the component.
3. Write to the MSMON_CFG_* registers for the instance of the monitor type.
4. Repeat step 3 to configure additional registers associated with the monitor instance.
5. Repeat steps 2 through 4 to configure additional monitor instances.
6. Release the exclusive access to the MSC's monitor configuration registers (for example, release the lock taken in step 1).

Repeat this procedure for each MSC.

Software must also take the lock to read any MSMON_* register, other than MSMON_CFG_MON_SEL, because reading involves first writing a monitor index into MSMON_CFG_MON_SEL and then reading an MSMON register.

The monitor configuration registers are all of the registers in *Memory-mapped monitoring configuration registers*. These registers have requirements similar to the MPAMCFG_* registers. The monitor configuration registers can have a separate lock or share the same lock as for the MPAMCFG_* registers. The selection register for monitors is MSMON_CFG_MON_SEL.

The configuration reading procedure of this section is also required to read the monitor and capture registers because these too are addressed by MSMON_CFG_MON_SEL.

There are two copies of MSMON_CFG_MON_SEL, one for Secure monitors that are accessed from the Secure address space and the other for Non-secure monitors that are accessed from the Non-secure address space. Because there are two copies, there can be separate locks for Secure MSMON_CFG_MON_SEL and for Non-secure MSMON_CFG_MON_SEL.

9.1.4 MPAM feature page

An MSC has an MPAM feature page in each of the supported address spaces. An MPAM feature page is a block of addresses that contains all of the MPAM MSC MMRs in that address space. Each MPAM feature page base address must be aligned to a 4 KB boundary.

Each MPAM feature page must be completely contained within a single 64 KB aligned block so that it may be placed within a single 64 KB page. Non-MPAM MMRs of the MSC are permitted within the 64 KB block if those MMRs are also to be trapped to a hypervisor.

Secure, Non-secure, Root and Realm address space

If the MSC supports the Secure address space (NS == 0), the Secure MPAM feature page must exist. The Non-secure MPAM feature page must always exist.

When FEAT_RME is implemented, the Root and Realm MPAM feature pages must exist. See *Four-space MSC*.

MMRs describing (IDRs) or controlling (MPAMCFG*) Secure PARTIDs are within the Secure MPAM feature page, and those describing or controlling Non-secure PARTIDs are within the Non-secure MPAM feature page. MMRs describing or controlling PARTIDs in the Root PARTID space are within the Root MPAM feature page, and those describing or controlling PARTIDs in the Realm PARTID space are within the Realm MPAM feature page.

MPAM MMRs only in the Secure address space

Certain MPAM MMRs are only present within the MPAM feature page when accessed via the Secure address space (NS = 0). MPAMF_SIDR is the only MMR accessible only via the Secure address space.

Read-only MPAM MMRs permitted to read the same or differently

Some of the read-only MPAM MMRs are permitted to have the same or different contents between the Secure, Non-secure, Root and Realm MPAM feature pages. This includes all of the MPAMF*IDR registers. If the information regarding Secure and Non-secure PARTIDs is the same in an MPAMF*IDR, then the register is permitted to have the same contents.

These registers are permitted to be shared if the same or banked if different in the two address spaces:

MPAMF_IDR	MPAMF_IMPL_IDR	MPAMF_CPOR_IDR
MPAMF_CCAP_IDR	MPAMF_MBW_IDR	MPAMF_PRI_IDR
MPAMF_PARTID_NRW_IDR	MPAMF_MSMON_IDR	MPAMF_CSUMON_IDR
MPAMF_MBWUMON_IDR		

MPAM MMRs that must have the same contents

Two registers must have the same contents between the Secure and Non-secure MPAM feature pages. These registers contain read-only values that must read as the same value in the two address spaces:

MPAMF_IIDR	MPAMF_AIDR
------------	------------

MPAM MMRs that must be separate registers for each address space

Most MPAM MMRs, such as the following, must be separate and have Secure, Non-secure, Root and Realm versions that are accessed through the corresponding Secure, Non-secure, Root and Realm MPAM feature pages:

MPAMF_ECR	MPAMCFG_PART_SEL	MSMON_CFG_MON_SEL
MPAMF_ESR	MPAMCFG_MBW_MAX	MSMON_CFG_CSU_CTL
	MPAMCFG_MBW_MIN	MSMON_CFG_CSU_FLT
MPAMCFG_CMAX	MPAMCFG_MBW_PBM	MSMON_CSU
MPAMCFG_CPBM	MPAMCFG_MBW_PROP	MSMON_CSU_CAPTURE
	MPAMCFG_MBW_WINWD	MSMON_CFG_MBWU_CTL

MPAMCFG_PRI	MSMON_CFG_MBWU_FLT
MPAMCFG_INTPARTID	MSMON_MBWU
	MSMON_MBWU_CAPTURE

Accesses to locations where there is no register in the address space of the access

Access to MPAM MMR address where there is no register in the address space of the access must be treated as reserved MPAM feature page locations according to *IMPLEMENTATION DEFINED memory-mapped registers and reserved feature page locations*, except for the MPAMCFG_MBW_PBM and MPAMCFG_CPBM as described in *Permitted truncation of an MPAM feature page*.

Permitted truncation of an MPAM feature page

An MPAM feature page may be shortened in only two cases:

- If MPAMCFG_MBW_PBM is not implemented ($MPAMF_IDR.HAS_MBW_PART == 0' || (MPAM_IDR.HAS_MBW_PART == 1 \&\& MPAM_MBW_IDR.HAS_PBM == 0)$), the maximum offset for the MPAM feature page is $0x01FFF$.
- If MPAMCFG_MBW_PBM is not implemented and MPAMCFG_CPBM is not implemented ($MPAMF_IDR.HAS_CPOR == 0$), the maximum offset for the MPAM feature page is $0x00FFF$.

9.1.5 Minimum required MPAM memory-mapped registers

If an MSC has any support for MPAM, the following registers are required:

- MPAMF_IDR.
- MPAMF_AIDR.
- MPAMF_IIDR.
- MPAMF_SIDR, if the Secure address space is supported.

If an MSC supports any resource controls, the following registers are also required:

- MPAMCFG_PART_SEL.

If an MSC supports any resource monitors, the following registers are also required:

- MPAMF_MSMON_IDR.
- MSMON_CFG_MON_SEL.

If an MSC can detect any errors, it must implement:

- MPAMF_ESR.
- MPAMF_ECR.

MSC MPAM MMRs not mentioned in this section are optional are expected to be implemented only when the resource control or monitor that the register supports is implemented.

See *Examples of partial MPAM implementations* for examples showing MPAMF_*IDR registers in implementations with few MPAM functions.

9.1.6 IMPLEMENTATION DEFINED memory-mapped registers and reserved feature page locations

IMPLEMENTATION DEFINED MPAM memory-mapped registers are permitted in the MPAM feature page at offsets equal to or greater than $0x3000$.

All locations in the MPAM feature page at offsets less than the maximum MPAM feature page offset defined in *Permitted truncation of an MPAM feature page* are reserved to the architecture. Within that address range:

- Reads and writes of unallocated locations are reserved accesses.

- Reads and writes of locations for registers that are not implemented are reserved accesses, including register locations for:
 - Optional MPAM MSC features that are not implemented.
 - ID registers for optional MPAM MSC features that are not implemented and indicated as not implemented in ID registers that are implemented.
- Locations that are beyond the implemented width of a register as given in the corresponding ID register but within the range of locations allocated by the architecture are reserved accesses.
- Reads of WO locations are reserved accesses.
- Writes to RO locations are reserved accesses.

The architecture requires reserved accesses to be implemented as RAZ/WI. However, software must not rely on this property as the behavior of reserved values might change in a future revision of the MPAM Extension architecture. Software must treat reserved accesses as RES0.

9.1.7 Examples of partial MPAM implementations

Most MSCs only implement a fraction of the full MSC MPAM architecture. This section gives examples of partial implementations, some of which have been achieved by partially removing MPAM. The RTL configuration examples are included to illustrate the MMR issues in partial MPAM implementations.

An MSC that has no partitioning or monitoring, only propagation

An MSC that does not implement any resource partitioning or monitor interfaces only requires a few MMRs:

- The minimum required MMRs, as specified in *Minimum required MPAM memory-mapped registers*, must be implemented with the `MPAMF_IDR`.{`PARTID_MAX`, `PMG_MAX`} fields indicating the maximum `PARTID` that can be propagated.
- All of the `HAS_*` and `NO_*` bits in `MPAMF_IDR` must be zero.
- `MPAMF_AIDR` must indicate MPAM v1.0 or MPAM v1.1.
- `MPAMF_IIDR` must identify the implementation.
- `MPAMF_SIDR` must indicate `PARTID_MAX` and `PMG_MAX` for Secure propagation.

No other registers are required.

An MSC when RTL configuration has removed a partitioning control or resource usage monitor

An MSC could be designed to have an RTL configuration option that removes a partitioning control or a resource usage monitor. If so, the `HAS_*` bits in each of the relevant `MPAMF_*IDR` registers must be configured to zero when the feature is removed.

An MSC when RTL configuration has removed all MPAM functionality

An MSC could be designed to have an RTL configuration option that removes all of the MPAM functionality. When all of MPAM is deconfigured:

- The minimum required MPAM registers must be present.
- `MPAMF_IDR`, `MPAMF_AIDR` and `MPAMF_SIDR` must all be zero.
- `MPAMF_IIDR` is permitted to be either all zero or to identify the IP.

————— Note —————

Software might still attempt to discover MPAM on this RTL configuration, so the minimum MPAM registers must be present to allow the lack of MPAM function to be discovered.

An MSC when RTL configuration removes a resource instance

An MSC could be designed to have an RTL configuration option that completely removes one or more resource instances. When a resource instance is removed, only the MPAMF_*IDR registers for the corresponding RIS values are changed. All of the ID registers corresponding to that RIS value have each of their RIS-specific fields set to zero. For more information on RIS-specific fields, see *Effects of MPAMCFG_PART_SEL.RIS on values read from other registers*.

9.2 Summary of memory-mapped registers

Table 9-1 lists the external MPAM registers in order of register offset.

Table 9-1 Index of external MPAM registers ordered by offset

Register	Offset	Length	Description, see:
MPAMF_IDR	0x0000	64	<i>MPAMF_IDR, MPAM Features Identification Register</i>
MPAMF_SIDR	0x0008	32	<i>MPAMF_SIDR, MPAM Features Secure Identification Register</i>
MPAMF_IIDR	0x0018	32	<i>MPAMF_IIDR, MPAM Implementation Identification Register</i>
MPAMF_AIDR	0x0020	32	<i>MPAMF_AIDR, MPAM Architecture Identification Register</i>
MPAMF_IMPL_IDR	0x0028	32	<i>MPAMF_IMPL_IDR, MPAM Implementation-Specific Partitioning Feature Identification Register</i>
MPAMF_CPOR_IDR	0x0030	32	<i>MPAMF_CPOR_IDR, MPAM Features Cache Portion Partitioning ID register</i>
MPAMF_CCAP_IDR	0x0038	32	<i>MPAMF_CCAP_IDR, MPAM Features Cache Capacity Partitioning ID register</i>
MPAMF_MBW_IDR	0x0040	32	<i>MPAMF_MBW_IDR, MPAM Memory Bandwidth Partitioning Identification Register</i>
MPAMF_PRI_IDR	0x0048	32	<i>MPAMF_PRI_IDR, MPAM Priority Partitioning Identification Register</i>
MPAMF_PARTID_NRW_IDR	0x0050	32	<i>MPAMF_PARTID_NRW_IDR, MPAM PARTID Narrowing ID register</i>
MPAMF_MSMON_IDR	0x0080	32	<i>MPAMF_MSMON_IDR, MPAM Resource Monitoring Identification Register</i>
MPAMF_CSUMON_IDR	0x0088	32	<i>MPAMF_CSUMON_IDR, MPAM Features Cache Storage Usage Monitoring ID register</i>
MPAMF_MBWUMON_IDR	0x0090	32	<i>MPAMF_MBWUMON_IDR, MPAM Features Memory Bandwidth Usage Monitoring ID register</i>
MPAMF_ERR_MSI_MPAM	0x00DC	32	<i>MPAMF_ERR_MSI_MPAM, MPAM Error MSI Write MPAM Information Register</i>
MPAMF_ERR_MSI_ADDR_L	0x00E0	32	<i>MPAMF_ERR_MSI_ADDR_L, MPAM Error MSI Low-part Address Register</i>
MPAMF_ERR_MSI_ADDR_H	0x00E4	32	<i>MPAMF_ERR_MSI_ADDR_H, MPAM Error MSI High-part Address Register</i>
MPAMF_ERR_MSI_DATA	0x00E8	32	<i>MPAMF_ERR_MSI_DATA, MPAM Error MSI Data Register</i>
MPAMF_ERR_MSI_ATTR	0x00EC	32	<i>MPAMF_ERR_MSI_ATTR, MPAM Error MSI Write Attributes Register</i>
MPAMF_ECR	0x00F0	32	<i>MPAMF_ECR, MPAM Error Control Register</i>
MPAMF_ESR	0x00F8	64	<i>MPAMF_ESR, MPAM Error Status Register</i>
MPAMCFG_PART_SEL	0x0100	32	<i>MPAMCFG_PART_SEL, MPAM Partition Configuration Selection Register</i>

Table 9-1 Index of external MPAM registers ordered by offset (continued)

Register	Offset	Length	Description, see:
MPAMCFG_CMAX	0x0108	32	<i>MPAMCFG_CMAX, MPAM Cache Maximum Capacity Partition Configuration Register</i>
MPAMCFG_CMIN	0x0110	32	<i>MPAMCFG_CMIN, MPAM Cache Minimum Capacity Partition Configuration Register</i>
MPAMCFG_CASSOC	0x0118	32	<i>MPAMCFG_CASSOC, MPAM Cache Maximum Associativity Partition Configuration Register</i>
MPAMCFG_MBW_MIN	0x0200	32	<i>MPAMCFG_MBW_MIN, MPAM Memory Bandwidth Minimum Partition Configuration Register</i>
MPAMCFG_MBW_MAX	0x0208	32	<i>MPAMCFG_MBW_MAX, MPAM Memory Bandwidth Maximum Partition Configuration Register</i>
MPAMCFG_MBW_WINWD	0x0220	32	<i>MPAMCFG_MBW_WINWD, MPAM Memory Bandwidth Partitioning Window Width Configuration Register</i>
MPAMCFG_EN	0x0300	32	<i>MPAMCFG_EN, MPAM Partition Configuration Enable Register</i>
MPAMCFG_DIS	0x0310	32	<i>MPAMCFG_DIS, MPAM Partition Configuration Disable Register</i>
MPAMCFG_EN_FLAGS	0x0320	32	<i>MPAMCFG_EN_FLAGS, MPAM Partition Configuration Enable Flags Register</i>
MPAMCFG_PRI	0x0400	32	<i>MPAMCFG_PRI, MPAM Priority Partition Configuration Register</i>
MPAMCFG_MBW_PROP	0x0500	32	<i>MPAMCFG_MBW_PROP, MPAM Memory Bandwidth Proportional Stride Partition Configuration Register</i>
MPAMCFG_INTPARTID	0x0600	32	<i>MPAMCFG_INTPARTID, MPAM Internal PARTID Narrowing Configuration Register</i>
MSMON_CFG_MON_SEL	0x0800	32	<i>MSMON_CFG_MON_SEL, MPAM Monitor Instance Selection Register</i>
MSMON_CAPT_EVNT	0x0808	32	<i>MSMON_CAPT_EVNT, MPAM Capture Event Generation Register</i>
MSMON_CFG_CSU_FLT	0x0810	32	<i>MSMON_CFG_CSU_FLT, MPAM Memory System Monitor Configure Cache Storage Usage Monitor Filter Register</i>
MSMON_CFG_CSU_CTL	0x0818	32	<i>MSMON_CFG_CSU_CTL, MPAM Memory System Monitor Configure Cache Storage Usage Monitor Control Register</i>
MSMON_CFG_MBWU_FLT	0x0820	32	<i>MSMON_CFG_MBWU_FLT, MPAM Memory System Monitor Configure Memory Bandwidth Usage Monitor Filter Register</i>
MSMON_CFG_MBWU_CTL	0x0828	32	<i>MSMON_CFG_MBWU_CTL, MPAM Memory System Monitor Configure Memory Bandwidth Usage Monitor Control Register</i>
MSMON_CSU	0x0840	32	<i>MSMON_CSU, MPAM Cache Storage Usage Monitor Register</i>
MSMON_CSU_CAPTURE	0x0848	32	<i>MSMON_CSU_CAPTURE, MPAM Cache Storage Usage Monitor Capture Register</i>
MSMON_CSU_OFSR	0x0858	32	<i>MSMON_CSU_OFSR, MPAM CSU Monitor Overflow Status Register</i>

Table 9-1 Index of external MPAM registers ordered by offset (continued)

Register	Offset	Length	Description, see:
MSMON_MBWU	0x0860	32	<i>MSMON_MBWU</i> , MPAM Memory Bandwidth Usage Monitor Register
MSMON_MBWU_CAPTURE	0x0868	32	<i>MSMON_MBWU_CAPTURE</i> , MPAM Memory Bandwidth Usage Monitor Capture Register
MSMON_MBWU_L	0x0880	64	<i>MSMON_MBWU_L</i> , MPAM Long Memory Bandwidth Usage Monitor Register
MSMON_MBWU_L_CAPTURE	0x0890	64	<i>MSMON_MBWU_L_CAPTURE</i> , MPAM Long Memory Bandwidth Usage Monitor Capture Register
MSMON_MBWU_OFSR	0x0898	32	<i>MSMON_MBWU_OFSR</i> , MPAM MBWU Monitor Overflow Status Register
MSMON_OFLOW_MSI_MPAM	0x08DC	32	<i>MSMON_OFLOW_MSI_MPAM</i> , MPAM Monitor Overflow MSI Write MPAM Information Register
MSMON_OFLOW_MSI_ADDR_L	0x08E0	32	<i>MSMON_OFLOW_MSI_ADDR_L</i> , MPAM Monitor Overflow MSI Low-part Address Register
MSMON_OFLOW_MSI_ADDR_H	0x08E4	32	<i>MSMON_OFLOW_MSI_ADDR_H</i> , MPAM Monitor Overflow MSI Write High-part Address Register
MSMON_OFLOW_MSI_DATA	0x08E8	32	<i>MSMON_OFLOW_MSI_DATA</i> , MPAM Monitor Overflow MSI Write Data Register
MSMON_OFLOW_MSI_ATTR	0x08EC	32	<i>MSMON_OFLOW_MSI_ATTR</i> , MPAM Monitor Overflow MSI Write Attributes Register
MSMON_OFLOW_SR	0x08F0	32	<i>MSMON_OFLOW_SR</i> , MPAM Monitor Overflow Status Register
MPAMCFG_CPBM<n>	0x1000	32	<i>MPAMCFG_CPBM<n></i> , MPAM Cache Portion Bitmap Partition Configuration Register; n = 0 - 1023
MPAMCFG_MBW_PBM<n>	0x2000	32	<i>MPAMCFG_MBW_PBM<n></i> , MPAM Bandwidth Portion Bitmap Partition Configuration Register; n = 0 - 127

9.3 Memory-mapped ID register description

This section lists the external ID registers.

9.3.1 MPAMF_AIDR, MPAM Architecture Identification Register

The MPAMF_AIDR characteristics are:

Purpose

Identifies the version of the MPAM architecture that this MSC implements.

Note: The following values are defined for bits [7:0]:

- 0x01 == MPAM architecture v0.1
- 0x10 == MPAM architecture v1.0
- 0x11 == MPAM architecture v1.1

Configurations

The power domain of MPAMF_AIDR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_AIDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_AIDR is a 32-bit register.

Field descriptions



Bits [31:8]

Reserved, RES0.

ArchMajorRev, bits [7:4]

Major revision of the MPAM architecture implemented by the MSC.

This table shows the only valid combinations of MPAM version numbers in an MSC. FORCE_NS functionality is only available in MPAM v0.1.

ArchMajorRev	ArchMinorRev	MPAMv	Available
0	0		None.
0	1	v0.1	MPAMv1.0 + MPAMv1.1 + FORCE_NS
1	0	v1.0	MPAMv1.0
1	1	v1.1	MPAMv1.0 + MPAMv1.1 - FORCE_NS

Use of MPAMv0.1 in MSCs is restricted to limited circumstances. The MSC must be able to initiate requests in the Secure address space which have MPAM PARTID forced to the Non-secure space with that forcing not controllable or observable by the software that configures the device for Secure requests. Please contact Arm before setting MPAMF_AIDR to report MPAMv0.1.

ArchMinorRev, bits [3:0]

Minor revision of the MPAM architecture implemented by the MSC.

See the table in the description of the ArchMajorRev field in this register.

Accessing the MPAMF_AIDR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_AIDR is read-only.

MPAMF_AIDR must be readable from the Secure, Non-secure, Root, and Realm MPAM feature pages.

MPAMF_AIDR must have the same contents in the Secure, Non-secure, Root, and Realm MPAM feature pages.

MPAMF_AIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0020	MPAMF_AIDR_s

Accesses to this interface are RO.

MPAMF_AIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0020	MPAMF_AIDR_ns

Accesses to this interface are RO.

MPAMF_AIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0020	MPAMF_AIDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_AIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0020	MPAMF_AIDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.2 MPAMF_CCAP_IDR, MPAM Features Cache Capacity Partitioning ID register

The MPAMF_CCAP_IDR characteristics are:

Purpose

Indicates the number of fractional bits in [MPAMCFG_CMAX.CMAX](#).

MPAMF_CCAP_IDR_s indicates the number of fractional bits in the Secure instance of [MPAMCFG_CMAX](#). MPAMF_CCAP_IDR_ns indicates the number of fractional bits in the Non-secure instance of [MPAMCFG_CMAX](#). MPAMF_CCAP_IDR_rt indicates the number of fractional bits in the Root cache capacity control settings register field, [MPAMCFG_CMAX.CMAX](#). MPAMF_CCAP_IDR_rl indicates the number of fractional bits in the Realm cache capacity control settings register field, [MPAMCFG_CMAX.CMAX](#).

When [MPAMF_IDR.HAS_RIS](#) is 1, some fields in this register give information for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#). The description of every field that is affected by [MPAMCFG_PART_SEL.RIS](#) has information within the field description.

Configurations

The power domain of MPAMF_CCAP_IDR is IMPLEMENTATION DEFINED.


This register is present only when [FEAT_MPAM](#) is implemented and [MPAMF_IDR.HAS_CCAP_PART](#) == 1. Otherwise, direct accesses to MPAMF_CCAP_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_CCAP_IDR is a 32-bit register.

Field descriptions



HAS_CMAX_SOFTLIM, bit [31]

When [FEAT_MPAMv0p1](#) is implemented or [FEAT_MPAMv1p1](#) is implemented:

Has soft limiting selection field in [MPAMCFG_CMAX](#).

0b0 If [MPAMCFG_CMAX](#) is implemented, it has no SOFTLIM field and the maximum capacity is controlled with a hard limit.

0b1 If [MPAMCFG_CMAX](#) is implemented, that register has a SOFTLIMIT field to select between hard or soft limiting to the CMAX parameter.

If RIS is implemented, this field indicates selectable limiting for the cache maximum capacity control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Otherwise:

Reserved, RES0.

NO_CMAX, bit [30]

When [FEAT_MPAMv0p1](#) is implemented or [FEAT_MPAMv1p1](#) is implemented:

Does not have CMAX partitioning.

0b0 [MPAMCFG_CMAX](#) is implemented.

0b1 [MPAMCFG_CMAX](#) is not implemented.

If RIS is implemented, this field indicates the absence of a cache maximum capacity partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Otherwise:

Reserved, RES0.

HAS_CMIN, bit [29]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Has cache minimum capacity partitioning.

0b0 [MPAMCFG_CMIN](#) is not implemented.

0b1 [MPAMCFG_CMIN](#) is implemented.

If RIS is implemented, this field indicates the presence of a cache minimum capacity partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Otherwise:

Reserved, RES0.

HAS_CASSOC, bit [28]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Has cache maximum associativity partitioning.

0b0 [MPAMCFG_CASSOC](#) is not implemented.

0b1 [MPAMCFG_CASSOC](#) is implemented.

If RIS is implemented, this field indicates the presence of a cache maximum associativity partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Otherwise:

Reserved, RES0.

Bits [27:13]

Reserved, RES0.

CASSOC_WD, bits [12:8]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Number of fractional bits implemented in the cache associativity partitioning control, [MPAMCFG_CASSOC.CASSOC](#), of this MSC. See [MPAMCFG_CASSOC](#).

If RIS is implemented, this field indicates the number of fractional bits in the cache capacity partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Otherwise:

Reserved, RES0.

Bits [7:6]

Reserved, RES0.

CMAX_WD, bits [5:0]

Number of fractional bits implemented in the cache capacity partitioning control, [MPAMCFG_CMAX.CMAX](#), of this device. See [MPAMCFG_CMAX](#).

This field must contain a value from 1 to 16, inclusive.

If RIS is implemented, this field indicates the number of fractional bits in the cache capacity partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Accessing the MPAMF_CCAP_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_CCAP_IDR is read-only.

MPAMF_CCAP_IDR must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

MPAMF_CCAP_IDR is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- MPAMF_CCAP_IDR_s is permitted to have either the same or different contents to MPAMF_CCAP_IDR_ns, MPAMF_CCAP_IDR_rt, or MPAMF_CCAP_IDR_rl.
- MPAMF_CCAP_IDR_ns is permitted to have either the same or different contents to MPAMF_CCAP_IDR_rt or MPAMF_CCAP_IDR_rl.
- MPAMF_CCAP_IDR_rt is permitted to have either the same or different contents to MPAMF_CCAP_IDR_rl.

There must be separate registers in the Secure (MPAMF_CCAP_IDR_s), Non-secure (MPAMF_CCAP_IDR_ns), Root (MPAMF_CCAP_IDR_rt), and Realm (MPAMF_CCAP_IDR_rl) MPAM feature pages.

When MPAMF_IDR.HAS_RIS is 1, MPAMF_CCAP_IDR shows the configuration of cache capacity partitioning for the cache resource instance selected by MPAMCFG_PART_SEL.RIS. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

MPAMF_CCAP_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0038	MPAMF_CCAP_IDR_s

Accesses to this interface are RO.

MPAMF_CCAP_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0038	MPAMF_CCAP_IDR_ns

Accesses to this interface are RO.

MPAMF_CCAP_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0038	MPAMF_CCAP_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_CCAP_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0038	MPAMF_CCAP_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.3 MPAMF_CPOR_IDR, MPAM Features Cache Portion Partitioning ID register

The MPAMF_CPOR_IDR characteristics are:

Purpose

Indicates the number of bits in [MPAMCFG_CPBM<n>](#).

MPAMF_CPOR_IDR_s indicates the number of bits in the Secure instance of [MPAMCFG_CPBM<n>](#). MPAMF_CPOR_IDR_ns indicates the number of bits in the Non-secure instance of [MPAMCFG_CPBM<n>](#). MPAMF_CPOR_IDR_rt indicates the number of bits in the Root instance of [MPAMCFG_CPBM<n>](#). MPAMF_CPOR_IDR_rl indicates the number of bits in the Realm instance of [MPAMCFG_CPBM<n>](#).

When [MPAMF_IDR.HAS_RIS](#) is 1, some fields in this register give information for the resource instance selector, [MPAMCFG_PART_SEL.RIS](#). The description of every field that is affected by [MPAMCFG_PART_SEL.RIS](#) has information within the field description.

Configurations

The power domain of MPAMF_CPOR_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and [MPAMF_IDR.HAS_CPOR_PART](#) == 1. Otherwise, direct accesses to MPAMF_CPOR_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_CPOR_IDR is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

CPBM_WD, bits [15:0]

Number of bits in the cache portion partitioning bit map of this device. See [MPAMCFG_CPBM<n>](#).

This field must contain a value from 1 to 32768, inclusive. Values greater than 32 require a group of 32-bit registers to access the CPBM, up to 1024 if CPBM_WD is the largest value.

If RIS is implemented, this field indicates the number bits in the cache portion bitmap for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

Accessing the MPAMF_CPOR_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_CPOR_IDR is read-only.

MPAMF_CPOR_IDR must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

MPAMF_CPOR_IDR is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- MPAMF_CPOR_IDR_s is permitted to have either the same or different contents to MPAMF_CPOR_IDR_ns, MPAMF_CPOR_IDR_rt, or MPAMF_CPOR_IDR_rl.

- MPAMF_CPOR_IDR_ns is permitted to have either the same or different contents to MPAMF_CPOR_IDR_rt or MPAMF_CPOR_IDR_rl.
- MPAMF_CPOR_IDR_rt is permitted to have either the same or different contents to MPAMF_CPOR_IDR_rl.

There must be separate registers in the Secure (MPAMF_CPOR_IDR_s), Non-secure (MPAMF_CPOR_IDR_ns), Root (MPAMF_CPOR_IDR_rt), and Realm (MPAMF_CPOR_IDR_rl) MPAM feature pages.

When MPAMF_IDR.HAS_RIS is 1, MPAMF_CPOR_IDR shows the configuration of cache portion partitioning for the cache resource instance selected by MPAMCFG_PART_SEL.RIS. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

MPAMF_CPOR_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0030	MPAMF_CPOR_IDR_s

Accesses to this interface are RO.

MPAMF_CPOR_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0030	MPAMF_CPOR_IDR_ns

Accesses to this interface are RO.

MPAMF_CPOR_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0030	MPAMF_CPOR_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_CPOR_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0030	MPAMF_CPOR_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.4 MPAMF_CSUMON_IDR, MPAM Features Cache Storage Usage Monitoring ID register

The MPAMF_CSUMON_IDR characteristics are:

Purpose

Indicates the number of cache storage usage monitor instances and other properties of the CSU monitoring.

MPAMF_CSUMON_IDR_s indicates the number and properties of Secure cache storage usage monitoring. MPAMF_CSUMON_IDR_ns indicates the number and properties of Non-secure cache storage usage monitoring. MPAMF_CSUMON_IDR_rt indicates the number and properties of Root cache storage usage monitoring. MPAMF_CSUMON_IDR_rl indicates the number and properties of Realm cache storage usage monitoring.

If MPAMF_IDR.HAS_RIS is 1, fields that mention RIS must reflect the properties of the resource instance currently selected by MPAMCFG_PART_SEL.RIS. Fields that do not mention RIS are constant across all resource instances.

Configurations

The power domain of MPAMF_CSUMON_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MSMON == 1 and MPAMF_MSMON_IDR.MSMON_CSU == 1. Otherwise, direct accesses to MPAMF_CSUMON_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_CSUMON_IDR is a 32-bit register.

Field descriptions



HAS_CAPTURE, bit [31]

The implementation supports copying an MSMON_CSU to the corresponding MSMON_CSU_CAPTURE on a capture event.

0b0 MSMON_CSU_CAPTURE is not implemented and there is no support for capture events in the CSU monitor.

0b1 The MSMON_CSU_CAPTURE register is implemented and the CSU monitor supports the capture event behavior.

If RIS is implemented, this field indicates that CSU monitor capture is implemented for the resource instance selected by MPAMCFG_PART_SEL.RIS.

CSU_RO, bit [30]

The implementation of MSMON_CSU is read-only.

0b0 MSMON_CSU is read/write.

0b1 MSMON_CSU is read-only.

If RIS is implemented, this field indicates that the MSMON_CSU monitor register is read-only for the resource instance selected by MPAMCFG_PART_SEL.RIS.

HAS_XCL, bit [29]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Has filtering to exclude clean data and implements the `MSMON_CFG_CSU_FLT.XCL` field.

0b0 `MSMON_CFG_CSU_FLT` does not implement the XCL field.

0b1 `MSMON_CFG_CSU_FLT` implements the XCL field to exclude counting data in the clean state in the monitor instance.

If RIS is implemented, this field indicates that the `MSMON_CFG_CSU_FLT.XCL` field is implemented in the CSU monitor instances for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

Bit [28]

Reserved, RES0.

HAS_OFLOW_LNKG, bit [27]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Supports `MSMON_CFG_CSU_CTL.OFLOW_LNKG` field to control how overflow on an instance affects other monitor instances in this MSC.

0b0 Does not support CSU overflow linkage.

0b1 Supports CSU overflow linkage and the `MSMON_CFG_CSU_CTL.OFLOW_LNKG` field.

If RIS is implemented, this field indicates that `MSMON_CFG_CSU_CTL.OFLOW_LNKG` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_OFSR, bit [26]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

The CSU monitor overflow status bitmap register, `MSMON_CSU_OFSR`, is implemented.

0b0 `MSMON_CSU_OFSR` register is not implemented.

0b1 `MSMON_CSU_OFSR` register is implemented.

If RIS is implemented, this field indicates that CSU monitor overflow status bitmap register is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_CEVNT_OFLW, bit [25]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Supports `MSMON_CFG_CSU_CTL.CEVNT_OFLW` field which can enable the CSU monitor instance to perform overflow behaviors on a capture event.

0b0 Does not support `MSMON_CFG_CSU_CTL.CEVNT_OFLW`.

0b1 Supports `MSMON_CFG_CSU_CTL.CEVNT_OFLW`.

If RIS is implemented, this field indicates that `MSMON_CFG_CSU_CTL.CEVNT_OFLW` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_OFLOW_CAPT, bit [24]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Supports `MSMON_CFG_CSU_CTL.OFLOW_CAPT` field which can enable the CSU monitor instance to capture the monitor on an overflow.

0b0 Does not support `MSMON_CFG_CSU_CTL.OFLOW_CAPT`.

0b1 Supports `MSMON_CFG_CSU_CTL.OFLOW_CAPT`.

If RIS is implemented, this field indicates that `MSMON_CFG_CSU_CTL.OFLOW_CAPT` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

Bits [23:16]

Reserved, RES0.

NUM_MON, bits [15:0]

The number of cache storage usage monitor instances implemented.

The largest `MSMON_CFG_MON_SEL.MON_SEL` value is `NUM_MON` minus 1.

If RIS is implemented, this field indicates the number of CSU monitor instances implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Accessing the MPAMF_CSUMON_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

`MPAMF_CSUMON_IDR` is read-only.

`MPAMF_CSUMON_IDR` must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

`MPAMF_CSUMON_IDR` is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- `MPAMF_CSUMON_IDR_s` is permitted to have either the same or different contents to `MPAMF_CSUMON_IDR_ns`, `MPAMF_CSUMON_IDR_rt`, or `MPAMF_CSUMON_IDR_rl`.
- `MPAMF_CSUMON_IDR_ns` is permitted to have either the same or different contents to `MPAMF_CSUMON_IDR_rt` or `MPAMF_CSUMON_IDR_rl`.
- `MPAMF_CSUMON_IDR_rt` is permitted to have either the same or different contents to `MPAMF_CSUMON_IDR_rl`.

There must be separate registers in the Secure (`MPAMF_CSUMON_IDR_s`), Non-secure (`MPAMF_CSUMON_IDR_ns`), Root (`MPAMF_CSUMON_IDR_rt`), and Realm (`MPAMF_CSUMON_IDR_rl`) MPAM feature pages.

When `MPAMF_IDR.HAS_RIS` is 1, `MPAMF_CSUMON_IDR` shows the configuration of cache storage usage monitoring for the cache resource instance selected by `MPAMCFG_PART_SEL.RIS`. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

Access to `MPAMF_CSUMON_IDR` is not affected by `MSMON_CFG_MON_SEL.RIS`.

`MPAMF_CSUMON_IDR` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_s</code>	0x0088	<code>MPAMF_CSUMON_IDR_s</code>

Accesses to this interface are RO.

MPAMF_CSUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0088	MPAMF_CSUMON_IDR_ns

Accesses to this interface are RO.

MPAMF_CSUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0088	MPAMF_CSUMON_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_CSUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0088	MPAMF_CSUMON_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.5 MPAMF_IDR, MPAM Features Identification Register

The MPAMF_IDR characteristics are:

Purpose

Indicates which memory partitioning and monitoring features are present on this MSC.

MPAMF_IDR_s indicates the MPAM features accessed from the Secure MPAM feature page.

MPAMF_IDR_ns indicates the MPAM features accessed from the Non-secure MPAM feature page.

MPAMF_IDR_rt indicates the MPAM features accessed from the Root MPAM feature page.

MPAMF_IDR_rl indicates the MPAM features accessed from the Realm MPAM feature page.

When MPAMF_IDR.HAS_RIS is 1, some fields in this register give information for the resource instance selected by MPAMCFG_PART_SEL.RIS. The description of every field that is affected by MPAMCFG_PART_SEL.RIS has that information within the field description.

Configurations

The power domain of MPAMF_IDR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_IDR are RES0.

MPAMF_IDR is 64-bit register when MPAM v0.1 or v1.1 is implemented.

Otherwise, MPAMF_IDR is a 32-bit register.

The power and reset domain of each MSC component is specific to that component.

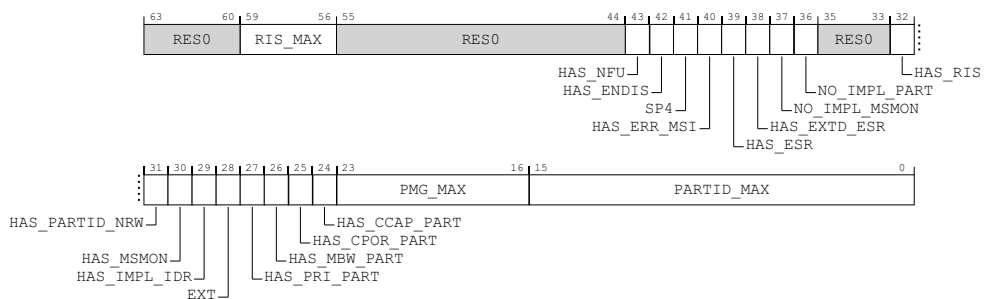
Attributes

MPAMF_IDR is a:

- 64-bit register when FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented
- 32-bit register otherwise

Field descriptions

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:



Bits [63:60]

Reserved, RES0.

RIS_MAX, bits [59:56]

When MPAMF_IDR.EXT == 1 and MPAMF_IDR.HAS_RIS == 1:

Maximum RIS value supported in MPAMCFG_PART_SEL. Must be 0b0000 if MPAMF_IDR.HAS_RIS == 0.

Otherwise:

Reserved, RES0.

Bits [55:44]

Reserved, RES0.

HAS_NFU, bit [43]

When FEAT_MPAMv1p1 is implemented or FEAT_MPAMv0p1 is implemented:

Has No Future Use field in MPAMCFG_DIS. Indicates that MPAMCFG_DIS.NFU is implemented.

0b0 MPAMCFG_DIS.NFU is not implemented. A PARTID disabled through access to MPAMCFG_DIS must preserve the control settings of the disabled PARTID.

0b1 Implements MPAMCFG_DIS.NFU. A PARTID disabled with NFU as 1 may have its control settings forgotten.

If MPAMF_IDR.HAS_ENDIS is 0b0, this field must also be 0b0.

This field must be the same in each instance of this register and for any value in MPAMCFG_PART_SEL.RIS.

Otherwise:

Reserved, RES0.

HAS_ENDIS, bit [42]

When FEAT_MPAMv1p1 is implemented or FEAT_MPAMv0p1 is implemented:

Has PARTID enable and disable. Indicates that this MSC supports PARTID disable and enable via MPAMCFG_DIS, MPAMCFG_EN and MPAMCFG_EN_FLAGS registers.

0b0 Does not support PARTID enable and disable functionality, and MPAMCFG_EN, MPAMCFG_DIS and MPAMCFG_EN_FLAGS registers are not implemented.

0b1 Supports PARTID enable and disable through the MPAMCFG_EN, MPAMCFG_DIS and MPAMCFG_EN_FLAGS registers.

All three registers must be implemented when this field is 1, MPAMCFG_EN, MPAMCFG_DIS, and MPAMCFG_EN_FLAGS.

This field must be the same in each instance of this register and for any value in MPAMCFG_PART_SEL.RIS.

Otherwise:

Reserved, RES0.

SP4, bit [41]

When FEAT_RME is implemented:

Indicates whether this MSC supports 4 PARTID spaces.

0b0 This MSC supports two PARTID spaces.

0b1 This MSC supports four PARTID spaces.

This field must read the same in each instance of this register and for any value in MPAMCFG_PART_SEL.RIS.

Otherwise:

Reserved, RES0.

HAS_ERR_MSI, bit [40]

When `MPAMF_IDR.EXT == 1`:

Has support for MSI writes to signal MPAM error interrupts. These registers are implemented: `MPAMF_ERR_MSI_ADDR_L`, `MPAMF_ERR_MSI_ADDR_H`, `MPAMF_ERR_MSI_ATTR`, `MPAMF_ERR_MSI_DATA`, and `MPAMF_ERR_MSI_MPAM`.

0b0 `MPAMF_ERR_MSI_ADDR_L`, `MPAMF_ERR_MSI_ADDR_H`, `MPAMF_ERR_MSI_ATTR`, `MPAMF_ERR_MSI_DATA`, and `MPAMF_ERR_MSI_MPAM` registers are not implemented.

0b1 `MPAMF_ERR_MSI_ADDR_L`, `MPAMF_ERR_MSI_ADDR_H`, `MPAMF_ERR_MSI_ATTR`, `MPAMF_ERR_MSI_DATA`, and `MPAMF_ERR_MSI_MPAM` are implemented and can be used to generate writes to signal error interrupts.

If `MPAMF_IDR.HAS_ESR` is 0, this bit must also be 0.

Otherwise:

Reserved, RES0.

HAS_ESR, bit [39]

When `MPAMF_IDR.EXT == 1`:

`MPAMF_ESR` is implemented.

0b0 `MPAMF_ESR`, `MPAMF_ECR`, and MPAM error handling are not implemented.

0b1 `MPAMF_ESR`, `MPAMF_ECR`, and MPAM error handling are implemented.

If an MSC cannot encounter any of the error conditions listed in [Errors in MSCs](#), both the `MPAMF_ESR` and `MPAMF_ECR` must be RAZ/WI.

Otherwise:

Reserved, RES0.

HAS_EXTD_ESR, bit [38]

When `MPAMF_IDR.EXT == 1`:

`MPAMF_ESR` is 64 bits.

0b0 `MPAMF_ESR` is 32 bits.

0b1 `MPAMF_ESR` is 64 bits.

When `MPAMF_IDR.HAS_RIS` and `MPAMF_IDR.HAS_ESR`, this field must be 1.

Otherwise:

Reserved, RES0.

NO_IMPL_MSMON, bit [37]

When `MPAMF_IDR.EXT == 1` and `MPAMF_IDR.HAS_IMPL_IDR == 1`:

`MPAMF_IMPL_IDR` defines no IMPLEMENTATION DEFINED resource monitors.

0b0 `MPAMF_IMPL_IDR` defines at least one IMPLEMENTATION DEFINED resource monitor.

0b1 `MPAMF_IMPL_IDR` does not define any IMPLEMENTATION DEFINED resource monitors.

If RIS is implemented, this field indicates the presence of IMPLEMENTATION DEFINED resource monitors described in `MPAMF_IMPL_IDR` for the selected resource instance.

Otherwise:

Reserved, RES0.

NO_IMPL_PART, bit [36]

When MPAMF_IDR.EXT == 1 and MPAMF_IDR.HAS_IMPL_IDR == 1:

MPAMF_IMPL_IDR defines no IMPLEMENTATION DEFINED resource controls.

0b0 MPAMF_IMPL_IDR defines at least one IMPLEMENTATION DEFINED resource control.

0b1 MPAMF_IMPL_IDR does not define any IMPLEMENTATION DEFINED resource controls.

If RIS is implemented, this field indicates the presence of IMPLEMENTATION DEFINED resource controls described in MPAMF_IMPL_IDR for the selected resource instance.

Otherwise:

Reserved, RES0.

Bits [35:33]

Reserved, RES0.

HAS_RIS, bit [32]

When MPAMF_IDR.EXT == 1:

Has resource instance selector. Indicates that MPAMCFG_PART_SEL contains the RIS field that selects a resource instance to control.

0b0 MPAMCFG_PART_SEL does not implement the MPAMCFG_PART_SEL.RIS field or multiple resource instance support.

0b1 MPAMCFG_PART_SEL implements the MPAMCFG_PART_SEL.RIS field and MPAM resource instance numbers up to and including MPAMF_IDR.RIS_MAX.

Otherwise:

Reserved, RES0.

HAS_PARTID_NRW, bit [31]

Has PARTID narrowing.

0b0 Does not have MPAMF_PARTID_NRW_IDR, MPAMCFG_INTPARTID, or intPARTID mapping support.

0b1 Supports the MPAMF_PARTID_NRW_IDR, MPAMCFG_INTPARTID registers.

HAS_MSMON, bit [30]

Has resource Monitors. Indicates whether this MSC has MPAM resource monitors.

0b0 Does not support MPAM resource monitoring by groups or MPAMF_MSMON_IDR.

0b1 Supports resource monitoring by matching a combination of PARTID and PMG. See MPAMF_MSMON_IDR.

HAS_IMPL_IDR, bit [29]

Has MPAMF_IMPL_IDR. Indicates whether this MSC has the IMPLEMENTATION SPECIFIC MPAM features register, MPAMF_IMPL_IDR.

0b0 Does not have MPAMF_IMPL_IDR.

0b1 Has MPAMF_IMPL_IDR.

EXT, bit [28]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Extended MPAMF_IDR.

0b0 MPAMF_IDR has no defined bits in [63:32]. The register is effectively 32 bits.

0b1 MPAMF_IDR has bits defined in [63:32]. The register is 64-bits.

Otherwise:

Reserved, RES0.

HAS_PRI_PART, bit [27]

Has Priority Partitioning. Indicates that MPAM priority partitioning is implemented and [MPAMF_PRI_IDR](#) exists.

0b0 Does not support priority partitioning or have [MPAMF_PRI_IDR](#).

0b1 Has priority partitioning and [MPAMF_PRI_IDR](#).

If RIS is implemented, this field indicates the presence of priority partitioning resource controls as described in [MPAMF_PRI_IDR](#) for the selected resource instance.

HAS_MBW_PART, bit [26]

Has Memory Bandwidth Partitioning. Indicates whether this MSC implements MPAM memory bandwidth partitioning and [MPAMF_MBW_IDR](#).

0b0 Does not support memory bandwidth partitioning or have [MPAMF_MBW_IDR](#) register.

0b1 Has [MPAMF_MBW_IDR](#) register.

If RIS is implemented, this field indicates the presence of memory bandwidth partitioning resource controls as described in [MPAMF_MBW_IDR](#) for the selected resource instance.

HAS_CPOR_PART, bit [25]

Has Cache Portion Partitioning. Indicates whether this MSC implements MPAM cache portion partitioning and [MPAMF_CPOR_IDR](#).

0b0 Does not support cache portion partitioning or have [MPAMF_CPOR_IDR](#) or [MPAMCFG_CPBM<n>](#) registers.

0b1 Has [MPAMF_CPOR_IDR](#) and [MPAMCFG_CPBM<n>](#) registers.

If RIS is implemented, this field indicates the presence of cache portion partitioning resource controls as described in [MPAMF_CPOR_IDR](#) for the selected resource instance.

HAS_CCAP_PART, bit [24]

Has Cache Capacity Partitioning. Indicates whether this MSC implements MPAM cache capacity partitioning and the [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.

0b0 Does not support cache capacity partitioning or have [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.

0b1 Has [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.

If RIS is implemented, this field indicates the presence of cache capacity partitioning resource controls as described in [MPAMF_CCAP_IDR](#) for the selected resource instance.

PMG_MAX, bits [23:16]

Maximum supported value of PMG.

The value of this field is permitted to vary between the instances of [MPAMF_IDR](#), each reporting the maximum supported PMG value in the PARTID space associated with that instance.

In [MPAMF_IDR_s](#), this field is permitted to report the maximum PMG value for the Non-secure PARTID space or for the Secure PARTID space. The maximum PMG value for the Secure PARTID space can be read from [MPAMF_SIDR.PMG_MAX](#).


PARTID_MAX, bits [15:0]

Maximum supported value of PARTID.

The value of this field is permitted to vary between the instances of [MPAMF_IDR](#), each reporting the maximum supported PARTID value in the PARTID space associated with that instance.

In [MPAMF_IDR_s](#), this field is permitted to report the maximum PARTID value for the Non-secure PARTID space or for the Secure PARTID space. The maximum PARTID value for the Secure PARTID space can be read from [MPAMF_SIDR.PARTID_MAX](#).

Otherwise:



HAS_PARTID_NRW, bit [31]

Has PARTID Narrowing.

- 0b0 Does not have [MPAMF_PARTID_NRW_IDR](#), [MPAMCFG_INTPARTID](#), or [intPARTID](#) mapping support.
- 0b1 Supports the [MPAMF_PARTID_NRW_IDR](#), [MPAMCFG_INTPARTID](#) registers.

HAS_MSMON, bit [30]

Has resource Monitors. Indicates whether this MSC has MPAM resource monitors.

- 0b0 Does not support MPAM resource monitoring by groups or [MPAMF_MSMON_IDR](#).
- 0b1 Supports resource monitoring by matching a combination of PARTID and PMG. See [MPAMF_MSMON_IDR](#).

HAS_IMPL_IDR, bit [29]

Has [MPAMF_IMPL_IDR](#). Indicates whether this MSC has the IMPLEMENTATION SPECIFIC MPAM features register, [MPAMF_IMPL_IDR](#).

- 0b0 Does not have [MPAMF_IMPL_IDR](#).
- 0b1 Has [MPAMF_IMPL_IDR](#).

EXT, bit [28]

When [FEAT_MPAMv0p1](#) is implemented or [FEAT_MPAMv1p1](#) is implemented:

Extended [MPAMF_IDR](#).

- 0b0 [MPAMF_IDR](#) has no defined bits in [63:32]. The register is effectively 32 bits.
- 0b1 [MPAMF_IDR](#) has bits defined in [63:32]. The register is 64-bits.

Otherwise:

Reserved, RES0.

HAS_PRI_PART, bit [27]

Has Priority Partitioning. Indicates whether this MSC implements MPAM priority partitioning and [MPAMF_PRI_IDR](#).

- 0b0 Does not support priority partitioning or have [MPAMF_PRI_IDR](#).
- 0b1 Has [MPAMF_PRI_IDR](#).

HAS_MBW_PART, bit [26]

Has Memory Bandwidth Partitioning. Indicates whether this MSC implements MPAM memory bandwidth partitioning and [MPAMF_MBW_IDR](#).

- 0b0 Does not support memory bandwidth partitioning or have [MPAMF_MBW_IDR](#) register.
- 0b1 Has [MPAMF_MBW_IDR](#) register.

HAS_CPOR_PART, bit [25]

Has Cache Portion Partitioning. Indicates whether this MSC implements MPAM cache portion partitioning and [MPAMF_CPOR_IDR](#).

- 0b0 Does not support cache portion partitioning or have [MPAMF_CPOR_IDR](#) or [MPAMCFG_CPBM<n>](#) registers.
- 0b1 Has [MPAMF_CPOR_IDR](#) and [MPAMCFG_CPBM<n>](#) registers.

HAS_CCAP_PART, bit [24]

Has Cache Capacity Partitioning. Indicates whether this MSC implements MPAM cache capacity partitioning and the [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.

- 0b0 Does not support cache capacity partitioning or have [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.
- 0b1 Has [MPAMF_CCAP_IDR](#) and [MPAMCFG_CMAX](#) registers.

PMG_MAX, bits [23:16]

Maximum supported value of PMG.

The value of this field is permitted to vary between the instances of [MPAMF_IDR](#), each reporting the maximum supported PMG value in the PARTID space associated with that instance.

In [MPAMF_IDR_s](#) this field is permitted to report the maximum PMG value for the Non-secure PARTID space or for the Secure PARTID space. The maximum PMG value for the Secure PARTID space can be read from [MPAMF_SIDR.PMG_MAX](#).

PARTID_MAX, bits [15:0]

Maximum supported value of PARTID.

The value of this field is permitted to vary between the instances of [MPAMF_IDR](#), each reporting the maximum supported PARTID value in the PARTID space associated with that instance.

In [MPAMF_IDR_s](#) this field is permitted to report the maximum PARTID value for the Non-secure PARTID space or for the Secure PARTID space. The maximum PARTID value for the Secure PARTID space can be read from [MPAMF_SIDR.PARTID_MAX](#).

Accessing the MPAMF_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

[MPAMF_IDR](#) is read-only.

[MPAMF_IDR](#) must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

[MPAMF_IDR](#) is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- [MPAMF_IDR_s](#) is permitted to have either the same or different contents to [MPAMF_IDR_ns](#), [MPAMF_IDR_rt](#), or [MPAMF_IDR_rl](#).
- [MPAMF_IDR_ns](#) is permitted to have either the same or different contents to [MPAMF_IDR_rt](#) or [MPAMF_IDR_rl](#).
- [MPAMF_IDR_rt](#) is permitted to have either the same or different contents to [MPAMF_IDR_rl](#).

There must be separate registers in the Secure ([MPAMF_IDR_s](#)), Non-secure ([MPAMF_IDR_ns](#)), Root ([MPAMF_IDR_rt](#)), and Realm ([MPAMF_IDR_rl](#)) MPAM feature pages.

When [MPAMF_IDR.HAS_RIS](#) is 1, [MPAMF_IDR](#) shows the configuration of MSC MPAM for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#). Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

MPAMF_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0000	MPAMF_IDR_s

Accesses to this interface are RO.

MPAMF_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0000	MPAMF_IDR_ns

Accesses to this interface are RO.

MPAMF_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0000	MPAMF_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0000	MPAMF_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.6 MPAMF_IIDR, MPAM Implementation Identification Register

The MPAMF_IIDR characteristics are:

Purpose

Uniquely identifies the MSC implementation by the combination of implementer, product ID, variant, and revision.

Configurations

The power domain of MPAMF_IIDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_IIDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_IIDR is a 32-bit register.

Field descriptions



ProductID, bits [31:20]

The MSC implementer as identified in the MPAMF_IIDR.Implementer field must assure each product has a unique ProductID from any other with the same Implementer value.

This field has an IMPLEMENTATION DEFINED value.

Access to this field is RO.

Variant, bits [19:16]

This field distinguishes product variants or major revisions of the product.

———— Note ————

Implementations of ProductID with differing software interfaces are expected to have different values in the MPAMF_IIDR.Variant field.

This field has an IMPLEMENTATION DEFINED value.

Access to this field is RO.

Revision, bits [15:12]

This field distinguishes minor revisions of the product.

———— Note ————

This field is intended to differentiate product revisions that are minor changes and are largely software compatible with previous revisions.

This field has an IMPLEMENTATION DEFINED value.

Access to this field is RO.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the MPAM MSC.

[11:8] must contain the JEP106 continuation code of the implementer.

[7] must always be 0.

[6:0] must contain the JEP106 identity code of the implementer.

For an Arm implementation, bits[11:0] are 0x43B.

Accessing the MPAMF_IIDR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_IIDR is read-only.

MPAMF_IIDR must be readable from the Secure, Non-secure, Root, and Realm MPAM feature pages.

MPAMF_IIDR must have the same contents in the Secure, Non-secure, Root, and Realm MPAM feature pages.

MPAMF_IIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0018	MPAMF_IIDR_s

Accesses to this interface are RO.

MPAMF_IIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0018	MPAMF_IIDR_ns

Accesses to this interface are RO.

MPAMF_IIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0018	MPAMF_IIDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_IIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0018	MPAMF_IIDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.7 MPAMF_IMPL_IDR, MPAM Implementation-Specific Partitioning Feature Identification Register

The MPAMF_IMPL_IDR characteristics are:

Purpose

Indicates the implementation-defined partitioning and monitoring features and parameters of the MSC.

MPAMF_IMPL_IDR_s indicates IMPLEMENTATION DEFINED partitioning and monitoring features accessed from the Secure MPAM feature page. MPAMF_IMPL_IDR_ns indicates those accessed from the Non-secure MPAM feature page. MPAMF_IMPL_IDR_rt indicates IMPLEMENTATION DEFINED partitioning and monitoring features accessed from the Root MPAM feature page. MPAMF_IMPL_IDR_rl indicates those accessed from the Realm MPAM feature page.

If MPAMF_IDR.HAS_RIS is 1, this register gives the implementation-specific features and parameters of the resource instance selected by MPAMCFG_PART_SEL.RIS for any features that are specific to the resource.

Configurations

The power domain of MPAMF_IMPL_IDR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_IMPL_IDR == 1. Otherwise, direct accesses to MPAMF_IMPL_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_IMPL_IDR is a 32-bit register.

Field descriptions



IMPLFEAT, bits [31:0]

All 32 bits of this register are available to be used as the implementer sees fit to indicate the presence of IMPLEMENTATION DEFINED MPAM features in this MSC and to give additional implementation-specific read-only information about the parameters of implementation-specific MPAM features to software.

If RIS is implemented, this register indicates the implementation-specific features and parameters of the resource instance selected by MPAMCFG_PART_SEL.RIS.

Accessing the MPAMF_IMPL_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_IMPL_IDR is read-only.

MPAMF_IMPL_IDR must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

MPAMF_IMPL_IDR is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- MPAMF_IMPL_IDR_s is permitted to have either the same or different contents to MPAMF_IMPL_IDR_ns, MPAMF_IMPL_IDR_rt, or MPAMF_IMPL_IDR_rl.
- MPAMF_IMPL_IDR_ns is permitted to have either the same or different contents to MPAMF_IMPL_IDR_rt or MPAMF_IMPL_IDR_rl.

- MPAMF_IMPL_IDR_rt is permitted to have either the same or different contents to MPAMF_IMPL_IDR_rl.

There must be separate registers in the Secure (MPAMF_IMPL_IDR_s), Non-secure (MPAMF_IMPL_IDR_ns), Root (MPAMF_IMPL_IDR_rt), and Realm (MPAMF_IMPL_IDR_rl) MPAM feature pages.

When MPAMF_IDR.HAS_RIS is 1, MPAMF_IMPL_IDR shows the configuration of implementation-specific features for the resource instance selected by MPAMCFG_PART_SEL.RIS. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

MPAMF_IMPL_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0028	MPAMF_IMPL_IDR_s

Accesses to this interface are RO.

MPAMF_IMPL_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0028	MPAMF_IMPL_IDR_ns

Accesses to this interface are RO.

MPAMF_IMPL_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0028	MPAMF_IMPL_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_IMPL_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0028	MPAMF_IMPL_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.8 MPAMF_MBW_IDR, MPAM Memory Bandwidth Partitioning Identification Register

The MPAMF_MBW_IDR characteristics are:

Purpose

Indicates which MPAM bandwidth partitioning features are present on this MSC.

MPAMF_MBW_IDR_s indicates bandwidth partitioning features accessed from the Secure MPAM feature page. MPAMF_MBW_IDR_ns indicates bandwidth partitioning features accessed from the Non-secure MPAM feature page. MPAMF_MBW_IDR_rt indicates bandwidth partitioning features accessed from the Root MPAM feature page. MPAMF_MBW_IDR_rl indicates bandwidth partitioning features accessed from the Realm MPAM feature page.

When MPAMF_IDR.HAS_RIS is 1, some fields in this register give information for the resource instance selected by MPAMCFG_PART_SEL.RIS. The description of every field that is affected by MPAMCFG_PART_SEL.RIS has that information within the field description.

Configurations

The power domain of MPAMF_MBW_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_MBW_PART == 1. Otherwise, direct accesses to MPAMF_MBW_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_MBW_IDR is a 32-bit register.

Field descriptions



Bits [31:29]

Reserved, RES0.

BWPBM_WD, bits [28:16]

Bandwidth portion bitmap width.

The number of bandwidth portion bits in the MPAMCFG_MBW_PBM<n> register array.

If MPAMF_MBW_IDR.HAS_PBM is 1, this field must contain a value from 1 to 4096, inclusive. Values greater than 32 require a group of 32-bit registers to access the BWPBM, up to 128 if BWPBM_WD is the largest value.

If MPAMF_MBW_IDR.HAS_PBM is 0, this field must be ignored by software.

If RIS is implemented, this field indicates the width of the memory bandwidth portion bitmap partitioning control for the resource instance selected by MPAMCFG_PART_SEL.RIS.

Bit [15]

Reserved, RES0.

WINDWR, bit [14]

Indicates the bandwidth accounting period register is writable.

0b0 The bandwidth accounting period is readable from MPAMCFG_MBW_WINWD which might be fixed or vary due to clock rate reconfiguration of the memory channel or memory controller.

0b1 The bandwidth accounting width is readable and writable per partition in [MPAMCFG_MBW_WINWD](#).

HAS_PROP, bit [13]

Indicates that this MSC implements proportional stride bandwidth partitioning and the [MPAMCFG_MBW_PROP](#) register can be accessed.

0b0 There is no memory bandwidth proportional stride control and the [MPAMCFG_MBW_PROP](#) register is RES0.

0b1 The proportional stride memory bandwidth partitioning scheme is supported and the [MPAMCFG_MBW_PROP](#) register can be accessed.

If RIS is implemented, this field indicates the presence of the memory bandwidth proportional stride partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

HAS_PBM, bit [12]

Indicates that bandwidth portion partitioning is implemented and the [MPAMCFG_MBW_PBM<n>](#) register array can be accessed.

0b0 There is no memory bandwidth portion control and the [MPAMCFG_MBW_PBM<n>](#) is RES0.

0b1 The memory bandwidth portion allocation scheme exists and the [MPAMCFG_MBW_PBM<n>](#) register can be accessed.

If RIS is implemented, this field indicates the presence of the memory bandwidth portion partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

HAS_MAX, bit [11]

Indicates that this MSC implements maximum bandwidth partitioning and the [MPAMCFG_MBW_MAX](#) register can be accessed.

0b0 There is no maximum memory bandwidth control and the [MPAMCFG_MBW_MAX](#) register is RES0.

0b1 The maximum memory bandwidth allocation scheme is supported and the [MPAMCFG_MBW_MAX](#) register can be accessed. Software can discover which limit behaviors are implemented by reading from [MPAMF_MBW_IDR.MAX_LIM](#), and can set the limit behavior by writing into [MPAMCFG_MBW_MAX.HARDLIM](#).

If RIS is implemented, this field indicates the presence of the maximum bandwidth partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

HAS_MIN, bit [10]

Indicates that this MSC implements minimum bandwidth partitioning and the [MPAMCFG_MBW_MIN](#) register can be accessed.

0b0 There is no minimum memory bandwidth control and the [MPAMCFG_MBW_MIN](#) register is RES0.

0b1 The minimum memory bandwidth allocation scheme is supported and the [MPAMCFG_MBW_MIN](#) register can be accessed.

If RIS is implemented, this field indicates the presence of the minimum bandwidth partitioning control for the resource instance selected by [MPAMCFG_PART_SEL.RIS](#).

MAX_LIM, bits [9:8]

When [MPAMF_MBW_IDR.HAS_MAX](#) == 1:

Implemented maximum-bandwidth limit partitioning behaviors.

0b00 Both soft limit and hard limit behaviors are implemented.

0b01 Soft limit behavior is implemented.

0b10 Hard limit behavior is implemented.

0b11 Reserved.

Software can set the limit behavior by writing into `MPAMCFG_MBW_MAX.HARDLIM`.

If RIS is implemented, this field indicates the presence of the minimum bandwidth partitioning control for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

Bits [7:6] Reserved, RES0.

BWA_WD, bits [5:0]

Number of implemented bits in the bandwidth allocation fields: MIN, MAX, and STRIDE. See `MPAMCFG_MBW_MIN`, `MPAMCFG_MBW_MAX`, and `MPAMCFG_MBW_PROP`.

In any of these bandwidth allocation fields exist, this field must have a value from 1 to 16, inclusive. Otherwise, it is permitted to be 0.

If RIS is implemented, this field indicates the number of implemented bits in the bandwidth allocation control fields for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Accessing the MPAMF_MBW_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

`MPAMF_MBW_IDR` is read-only.

`MPAMF_MBW_IDR` must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

`MPAMF_MBW_IDR` is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- `MPAMF_MBW_IDR_s` is permitted to have either the same or different contents to `MPAMF_MBW_IDR_ns`, `MPAMF_MBW_IDR_rt`, or `MPAMF_MBW_IDR_rl`.
- `MPAMF_MBW_IDR_ns` is permitted to have either the same or different contents to `MPAMF_MBW_IDR_rt` or `MPAMF_MBW_IDR_rl`.
- `MPAMF_MBW_IDR_rt` is permitted to have either the same or different contents to `MPAMF_MBW_IDR_rl`.

There must be separate registers in the Secure (`MPAMF_MBW_IDR_s`), Non-secure (`MPAMF_MBW_IDR_ns`), Root (`MPAMF_MBW_IDR_rt`), and Realm (`MPAMF_MBW_IDR_rl`) MPAM feature pages.

When `MPAMF_IDR.HAS_RIS` is 1, `MPAMF_MBW_IDR` shows the configuration of memory bandwidth partitioning for the bandwidth resource instance selected by `MPAMCFG_PART_SEL.RIS`. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

`MPAMF_MBW_IDR` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_s</code>	<code>0x0040</code>	<code>MPAMF_MBW_IDR_s</code>

Accesses to this interface are RO.

`MPAMF_MBW_IDR` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_ns</code>	<code>0x0040</code>	<code>MPAMF_MBW_IDR_ns</code>

Accesses to this interface are RO.

MPAMF_MBW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0040	MPAMF_MBW_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_MBW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0040	MPAMF_MBW_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.9 MPAMF_MBWUMON_IDR, MPAM Features Memory Bandwidth Usage Monitoring ID register

The MPAMF_MBWUMON_IDR characteristics are:

Purpose

Indicates the number of memory bandwidth usage monitor instances implemented. This register also indicates several properties of MBWU monitoring, including whether the implementation supports capture, scaling, or long counters.

MPAMF_MBWUMON_IDR_s indicates the number of Secure memory bandwidth usage monitor instances. MPAMF_MBWUMON_IDR_ns indicates the number of Non-secure memory bandwidth usage monitor instances. MPAMF_MBWUMON_IDR_rt indicates the number of Root memory bandwidth usage monitor instances. MPAMF_MBWUMON_IDR_rl indicates the number of Realm memory bandwidth usage monitor instances.

If MPAMF_IDR.HAS_RIS is 1, fields that mention RIS must reflect the properties of the resource instance currently selected by MPAMCFG_PART_SEL.RIS. Fields that do not mention RIS are constant across all resource instances.

Configurations

The power domain of MPAMF_MBWUMON_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MSMON == 1 and MPAMF_MSMON_IDR.MSMON_MBWU == 1. Otherwise, direct accesses to MPAMF_MBWUMON_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_MBWUMON_IDR is a 32-bit register.

Field descriptions



HAS_CAPTURE, bit [31]

The implementation supports copying an MSMON_MBWU to the corresponding MSMON_MBWU_CAPTURE on a capture event.

- 0b0 MSMON_MBWU_CAPTURE is not implemented and there is no support for capture events in the MBWU monitor.
- 0b1 The MSMON_MBWU_CAPTURE register is implemented and the MBWU monitor supports the capture event behavior.

If RIS is implemented, this field indicates that MBWU monitor capture is implemented for the resource instance selected by MPAMCFG_PART_SEL.RIS.

If MPAMF_MBWUMON_IDR.HAS_LONG is 1, this also indicates that MSMON_MBWU_L_CAPTURE is implemented.

HAS_LONG, bit [30]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Indicates whether MSMON_MBWU_L is implemented.

If HAS_CAPTURE is 1, indicates whether MSMON_MBWU_L_CAPTURE is implemented.

- 0b0 Does not implement MSMON_MBWU_L or MSMON_MBWU_L_CAPTURE.

0b1 Implements `MSMON_MBWU_L`. If `HAS_CAPTURE == 1`, `MSMON_MBWU_L_CAPTURE` is also implemented.

If RIS is implemented, this field indicates that the long MBWU monitor is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

If `MPAMF_MBWUMON_IDR.HAS_CAPTURE` is 1, this also indicates that `MSMON_MBWU_L_CAPTURE` is implemented.

Otherwise:

Reserved, RES0.

LWD, bit [29]

When `FEAT_MPAMv0p1` is implemented or `FEAT_MPAMv1p1` is implemented:

Long register VALUE width.

If `MPAMF_MBWUMON_IDR.HAS_LONG` is 0, `MPAMF_MBWUMON_IDR.LWD` must also be 0.

0b0 If `MPAMF_MBWUMON_IDR.HAS_LONG` is 1, `MSMON_MBWU_L` has 44-bit VALUE field in bits [43:0]. Bits [62:44] are RES0. If `HAS_LONG` is 1 and `MPAMF_MBWUMON_IDR.HAS_CAPTURE` is 1, `MSMON_MBWU_L_CAPTURE` also has 44-bit VALUE field in bits [43:0].

0b1 `MSMON_MBWU_L` has 63-bit VALUE field in bits [62:0]. If `MPAMF_MBWUMON_IDR.HAS_CAPTURE == 1`, `MSMON_MBWU_L_CAPTURE` also has 63-bit VALUE field in bits [62:0].

If RIS is implemented, this field indicates the length of the `MSMON_MBWU_L.VALUE` field implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_RWBW, bit [28]

When `FEAT_MPAMv0p1` is implemented or `FEAT_MPAMv1p1` is implemented:

Read/write bandwidth selection is implemented in `MSMON_CFG_MBWU_FLT`.

0b0 Read/write bandwidth selection is not implemented.

0b1 Read/write bandwidth selection is implemented.

If RIS is implemented, this field indicates whether read/write bandwidth collection selection is available in `MSMON_CFG_MBWU_FLT` for resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_OFLOW_LNKG, bit [27]

When `FEAT_MPAMv0p1` is implemented or `FEAT_MPAMv1p1` is implemented:

Supports `MSMON_CFG_MBWU_CTL.OFLOW_LNKG` field to control how overflow on an instance affects other monitor instances in this MSC.

0b0 Does not support MBWU overflow linkage.

0b1 Supports MBWU overflow linkage and the `MSMON_CFG_MBWU_CTL.OFLOW_LNKG` field.

If RIS is implemented, this field indicates that `MSMON_CFG_MBWU_CTL.OFLOW_LNKG` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_OFSR, bit [26]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

The MBWU monitor overflow status bitmap register, `MSMON_MBWU_OFSR`, is implemented.

0b0 `MSMON_MBWU_OFSR` register is not implemented.

0b1 `MSMON_MBWU_OFSR` register is implemented.

If RIS is implemented, this field indicates that MBWU monitor overflow status bitmap register is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

HAS_CEVNT_OFLW, bit [25]

Supports `MSMON_CFG_MBWU_CTL.CEVNT_OFLW` field which can enable the MBWU monitor instance to perform overflow behaviors on a capture event.

0b0 Does not support `MSMON_CFG_MBWU_CTL.CEVNT_OFLW`.

0b1 Supports `MSMON_CFG_MBWU_CTL.CEVNT_OFLW`.

If RIS is implemented, this field indicates that `MSMON_CFG_MBWU_CTL.CEVNT_OFLW` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

HAS_OFLOW_CAPT, bit [24]

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:

Supports `MSMON_CFG_MBWU_CTL.OFLOW_CAPT` field which can enable the MBWU monitor instance to capture the monitor on an overflow.

0b0 Does not support `MSMON_CFG_MBWU_CTL.OFLOW_CAPT`.

0b1 Supports `MSMON_CFG_MBWU_CTL.OFLOW_CAPT`.

If RIS is implemented, this field indicates that `MSMON_CFG_MBWU_CTL.OFLOW_CAPT` is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Otherwise:

Reserved, RES0.

Bits [23:21]

Reserved, RES0.

SCALE, bits [20:16]

Scaling of `MSMON_MBWU.VALUE` in bits. If scaling is enabled by `MSMON_CFG_MBWU_CTL.SCLEN`, the byte count in the `VALUE` field has been shifted by `SCALE` bits to the right.

0b00000 Scaling is not implemented.

0bxxxxx Other values are right shift count when scaling is enabled.

If RIS is implemented, this field indicates the scale value for `MSMON_MBWU.VALUE` field for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

NUM_MON, bits [15:0]

The number of memory bandwidth usage monitor instances implemented. The largest monitor instance selector, `MSMON_CFG_MON_SEL.MON_SEL`, is `NUM_MON` minus 1.

If RIS is implemented, this field indicates the number of MBWU monitor instances for `MSMON_MBWU.VALUE` field for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Accessing the MPAMF_MBWUMON_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_MBWUMON_IDR is read-only.

MPAMF_MBWUMON_IDR must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

MPAMF_MBWUMON_IDR is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- MPAMF_MBWUMON_IDR_s is permitted to have either the same or different contents to MPAMF_MBWUMON_IDR_ns, MPAMF_MBWUMON_IDR_rt, or MPAMF_MBWUMON_IDR_rl.
- MPAMF_MBWUMON_IDR_ns is permitted to have either the same or different contents to MPAMF_MBWUMON_IDR_rt or MPAMF_MBWUMON_IDR_rl.
- MPAMF_MBWUMON_IDR_rt is permitted to have either the same or different contents to MPAMF_MBWUMON_IDR_rl.

There must be separate registers in the Secure (MPAMF_MBWUMON_IDR_s), Non-secure (MPAMF_MBWUMON_IDR_ns), Root (MPAMF_MBWUMON_IDR_rt), and Realm (MPAMF_MBWUMON_IDR_rl) MPAM feature pages.

When MPAMF_IDR.HAS_RIS is 1, MPAMF_MBWUMON_IDR shows the configuration of memory bandwidth monitoring for the bandwidth resource instance selected by MPAMCFG_PART_SEL.RIS. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

Access to MPAMF_MBWUMON_IDR is not affected by MSMON_CFG_MON_SEL.RIS.

MPAMF_MBWUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0090	MPAMF_MBWUMON_IDR_s

Accesses to this interface are RO.

MPAMF_MBWUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0090	MPAMF_MBWUMON_IDR_ns

Accesses to this interface are RO.

MPAMF_MBWUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0090	MPAMF_MBWUMON_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_MBWUMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0090	MPAMF_MBWUMON_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.10 MPAMF_MSMON_IDR, MPAM Resource Monitoring Identification Register

The MPAMF_MSMON_IDR characteristics are:

Purpose

Indicates which MPAM monitoring features are present on this MSC.

MPAMF_MSMON_IDR_s indicates Secure monitoring features. MPAMF_MSMON_IDR_ns indicates Non-secure monitoring features. MPAMF_MSMON_IDR_rt indicates Root monitoring features. MPAMF_MSMON_IDR_rl indicates Realm monitoring features.

If MPAMF_IDR.HAS_RIS is 1, fields that mention RIS must reflect the properties of the resource instance currently selected by MPAMCFG_PART_SEL.RIS. Fields that do not mention RIS are constant across all resource instances.

Configurations

The power domain of MPAMF_MSMON_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_MSMON == 1. Otherwise, direct accesses to MPAMF_MSMON_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_MSMON_IDR is a 32-bit register.

Field descriptions



HAS_LOCAL_CAPT_EVNT, bit [31]

Has local capture event generator. Indicates whether this MSC has the MPAM local capture event generator and the MSMON_CAPT_EVNT register.

- 0b0 Does not support MPAM local capture event generator or MSMON_CAPT_EVNT.
- 0b1 Supports the MPAM local capture event generator and the MSMON_CAPT_EVNT register.

NO_HW_OFLW_INTR, bit [30]

When FEAT_MPAMv1p1 is implemented:

Does not have hardwired MPAM monitor overflow interrupt.

- 0b0 Supports generating a hardwired interrupt to signal MPAM monitor overflow.
- 0b1 No support for a hardwired interrupt to signal MPAM monitor overflow.

If this field is 0, the MSC supports generating a hardwired interrupt for monitor overflow events.

If this field is 0 and the HAS_OFLW_MSI field in this register is 1, the MSC supports generating both hardwired interrupts and MSI writes to signal interrupts.

Otherwise:

Reserved, RES0.

HAS_OFLW_MSI, bit [29]

When FEAT_MPAMv1p1 is implemented:

Has support for MSI writes to signal MPAM monitor overflow interrupts. These registers are implemented: `MSMON_OFLOW_MSI_ADDR_L`, `MSMON_OFLOW_MSI_ADDR_H`, `MSMON_OFLOW_MSI_ATTR`, `MSMON_OFLOW_MSI_DATA` and `MSMON_OFLOW_MSI_MPAM`.

0b0 `MSMON_OFLOW_MSI_ADDR_L`, `MSMON_OFLOW_MSI_ADDR_H`, `MSMON_OFLOW_MSI_ATTR`, `MSMON_OFLOW_MSI_DATA` and `MSMON_OFLOW_MSI_MPAM` registers are not implemented.

0b1 `MSMON_OFLOW_MSI_ADDR_L`, `MSMON_OFLOW_MSI_ADDR_H`, `MSMON_OFLOW_MSI_ATTR`, `MSMON_OFLOW_MSI_DATA` and `MSMON_OFLOW_MSI_ATTR` are implemented and can be used to generate writes to signal MPAM monitor overflow interrupts.

If `MPAMF_MSMON_IDR.NO_HW_OFLW_INTR` is 1 and this bit is 0, this MSC does not support monitor overflow interrupts.

Otherwise:

Reserved, RES0.

HAS_OFLOW_SR, bit [28]

When FEAT_MPAMv1p1 is implemented:

Has MPAM monitor overflow status register `MSMON_OFLOW_SR`.

0b0 Does not have `MSMON_OFLOW_SR`.

0b1 Supports `MSMON_OFLOW_SR`.

Otherwise:

Reserved, RES0.

Bits [27:18]

Reserved, RES0.

MSMON_MBWU, bit [17]

Memory bandwidth usage monitoring. Indicates whether MPAM monitoring for Memory Bandwidth Usage by PARTID and PMG is implemented and whether the following bandwidth usage registers are accessible:

- `MPAMF_MBWUMON_IDR`, `MSMON_CFG_MBWU_CTL`, `MSMON_CFG_MBWU_FLT`, `MSMON_MBWU`.
- The optional `MSMON_MBWU_CAPTURE`.
- If MPAM v0.1 or MPAM v1.1 is implemented, the optional `MSMON_MBWU_L` and the optional `MSMON_MBWU_L_CAPTURE`.

0b0 Does not have monitoring for memory bandwidth usage and does not use the bandwidth usage registers.

0b1 Has monitoring of memory bandwidth usage and uses the bandwidth usage registers.

If RIS is implemented, this field indicates that memory bandwidth usage monitoring is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS` as described in `MPAMF_MBWUMON_IDR`.

MSMON_CSU, bit [16]

Cache storage usage monitoring. Indicates whether MPAM monitoring of cache storage usage by PARTID and PMG is implemented and the following registers are accessible:

- `MPAMF_CSUMON_IDR`, `MSMON_CFG_CSU_CTL`, `MSMON_CFG_CSU_FLT`, `MSMON_CSU`.

- The optional `MSMON_CSU_CAPTURE`.
- 0b0 Does not have monitoring for cache storage usage or the `MPAMF_CSUMON_IDR`, `MSMON_CFG_CSU_CTL`, `MSMON_CFG_CSU_FLT`, `MSMON_CSU` or `MSMON_CSU_CAPTURE` registers.
- 0b1 Has monitoring of cache storage usage and the `MPAMF_CSUMON_IDR`, `MSMON_CFG_CSU_CTL`, `MSMON_CFG_CSU_FLT`, `MSMON_CSU` and optional `MSMON_CSU_CAPTURE` registers.

If RIS is implemented, this field indicates that cache storage usage monitoring is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS` as described in `MPAMF_CSUMON_IDR`.

Bits [15:0]

Reserved, RES0.

Accessing the `MPAMF_MSMON_IDR`:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

`MPAMF_MSMON_IDR` is read-only.

`MPAMF_MSMON_IDR` must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

`MPAMF_MSMON_IDR` is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- `MPAMF_MSMON_IDR_s` is permitted to have either the same or different contents to `MPAMF_MSMON_IDR_ns`, `MPAMF_MSMON_IDR_rt`, or `MPAMF_MSMON_IDR_rl`.
- `MPAMF_MSMON_IDR_ns` is permitted to have either the same or different contents to `MPAMF_MSMON_IDR_rt` or `MPAMF_MSMON_IDR_rl`.
- `MPAMF_MSMON_IDR_rt` is permitted to have either the same or different contents to `MPAMF_MSMON_IDR_rl`.

There must be separate registers in the Secure (`MPAMF_MSMON_IDR_s`), Non-secure (`MPAMF_MSMON_IDR_ns`), Root (`MPAMF_MSMON_IDR_rt`), and Realm (`MPAMF_MSMON_IDR_rl`) MPAM feature pages.

When `MPAMF_IDR.HAS_RIS` is 1, `MPAMF_MSMON_IDR` shows the configuration of memory system monitoring for the resource instance selected by `MPAMCFG_PART_SEL.RIS`. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

Access to `MPAMF_MSMON_IDR` is not affected by `MSMON_CFG_MON_SEL.RIS`.

`MPAMF_MSMON_IDR` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_s</code>	0x0080	<code>MPAMF_MSMON_IDR_s</code>

Accesses to this interface are RO.

`MPAMF_MSMON_IDR` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_ns</code>	0x0080	<code>MPAMF_MSMON_IDR_ns</code>

Accesses to this interface are RO.

MPAMF_MSMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0080	MPAMF_MSMON_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_MSMON_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0080	MPAMF_MSMON_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.11 MPAMF_PARTID_NRW_IDR, MPAM PARTID Narrowing ID register

The MPAMF_PARTID_NRW_IDR characteristics are:

Purpose

Indicates the largest internal PARTID for this MSC.

MPAMF_PARTID_NRW_IDR_s indicates the largest Secure internal PARTID.

MPAMF_PARTID_NRW_IDR_ns indicates the largest Non-secure internal PARTID.

When FEAT_RME is implemented: MPAMF_PARTID_NRW_rt indicates the largest Root internal PARTID. MPAMF_PARTID_NRW_rl indicates the largest Realm internal PARTID.

PARTID narrowing is global to the MSC and does not vary by resource instance.

Configurations

The power domain of MPAMF_PARTID_NRW_IDR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_PARTID_NRW == 1. Otherwise, direct accesses to MPAMF_PARTID_NRW_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_PARTID_NRW_IDR is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

INTPARTID_MAX, bits [15:0]

The largest intPARTID supported in this MSC.

Accessing the MPAMF_PARTID_NRW_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

MPAMF_PARTID_NRW_IDR is read-only.

MPAMF_PARTID_NRW_IDR must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

MPAMF_PARTID_NRW_IDR is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- MPAMF_PARTID_NRW_IDR_s is permitted to have either the same or different contents to MPAMF_PARTID_NRW_IDR_ns, MPAMF_PARTID_NRW_IDR_rt, or MPAMF_PARTID_NRW_IDR_rl.
- MPAMF_PARTID_NRW_IDR_ns is permitted to have either the same or different contents to MPAMF_PARTID_NRW_IDR_rt or MPAMF_PARTID_NRW_IDR_rl.
- MPAMF_PARTID_NRW_IDR_rt is permitted to have either the same or different contents to MPAMF_PARTID_NRW_IDR_rl.

There must be separate registers in the Secure (MPAMF_PARTID_NRW_IDR_s), Non-secure (MPAMF_PARTID_NRW_IDR_ns), Root (MPAMF_PARTID_NRW_IDR_rt), and Realm (MPAMF_PARTID_NRW_IDR_rl) MPAM feature pages.

MPAMF_PARTID_NRW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0050	MPAMF_PARTID_NRW_IDR_s

Accesses to this interface are RO.

MPAMF_PARTID_NRW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0050	MPAMF_PARTID_NRW_IDR_ns

Accesses to this interface are RO.

MPAMF_PARTID_NRW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0050	MPAMF_PARTID_NRW_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_PARTID_NRW_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0050	MPAMF_PARTID_NRW_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.12 MPAMF_PRI_IDR, MPAM Priority Partitioning Identification Register

The MPAMF_PRI_IDR characteristics are:

Purpose

Indicates which MPAM priority partitioning features are present on this MSC.

MPAMF_PRI_IDR_s indicates priority partitioning features accessed from the Secure MPAM feature page. MPAMF_PRI_IDR_ns indicates priority partitioning features accessed from the Non-secure MPAM feature page. MPAMF_PRI_IDR_rt indicates priority partitioning features accessed from the Root MPAM feature page. MPAMF_PRI_IDR_rl indicates priority partitioning features accessed from the Realm MPAM feature page.

When MPAMF_IDR.HAS_RIS is 1, some fields in this register give information for the resource instance selected by MPAMCFG_PART_SEL.RIS. The description of every field that is affected by MPAMCFG_PART_SEL.RIS has that information within the field description.

Configurations

The power domain of MPAMF_PRI_IDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_PRI_PART == 1. Otherwise, direct accesses to MPAMF_PRI_IDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_PRI_IDR is a 32-bit register.

Field descriptions



Bits [31:26]

Reserved, RES0.

DSPRI_WD, bits [25:20]

Number of implemented bits in the downstream priority field (DSPRI) of MPAMCFG_PRI.

If HAS_DSPRI == 1, this field must contain a value from 1 to 16, inclusive.

If HAS_DSPRI == 0, this field must be 0.

If RIS is implemented, this field indicates the number of downstream priority bits for the resource instance selected by MPAMCFG_PART_SEL.RIS.

Bits [19:18]

Reserved, RES0.

DSPRI_0_IS_LOW, bit [17]

Indicates whether 0 in MPAMCFG_PRI.DSPRI is the lowest or the highest downstream priority.

0b0 In the MPAMCFG_PRI.DSPRI field, a value of 0 means the highest priority.

0b1 In the MPAMCFG_PRI.DSPRI field, a value of 0 means the lowest priority.

If RIS is implemented, this field indicates that 0 is the lowest downstream priority for the resource instance selected by MPAMCFG_PART_SEL.RIS.

HAS_DSPRI, bit [16]

Indicates that the `MPAMCFG_PRI` register implements the DSPRI field.

- 0b0 This MSC supports priority partitioning, but does not implement a downstream priority (DSPRI) field in the `MPAMCFG_PRI` register.
- 0b1 This MSC supports downstream priority partitioning and implements the downstream priority (DSPRI) field in the `MPAMCFG_PRI` register.

If RIS is implemented, this field indicates that downstream priority is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Bits [15:10]

Reserved, RES0.

INTPRI_WD, bits [9:4]

Number of implemented bits in the internal priority field (INTPRI) in the `MPAMCFG_PRI` register.

If `HAS_INTPRI == 1`, this field must contain a value from 1 to 16, inclusive.

If `HAS_INTPRI == 0`, this field must be 0.

If RIS is implemented, this field indicates the number of internal priority bits for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Bits [3:2]

Reserved, RES0.

INTPRI_0_IS_LOW, bit [1]

Indicates whether 0 in `MPAMCFG_PRI.INTPRI` is the lowest or the highest internal priority.

- 0b0 In the `MPAMCFG_PRI.INTPRI` field, a value of 0 means the highest priority.
- 0b1 In the `MPAMCFG_PRI.INTPRI` field, a value of 0 means the lowest priority.

If RIS is implemented, this field indicates that 0 is the lowest internal priority for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

HAS_INTPRI, bit [0]

Indicates that this MSC implements the INTPRI field in the `MPAMCFG_PRI` register.

- 0b0 This MSC supports priority partitioning, but does not implement the internal priority (INTPRI) field in the `MPAMCFG_PRI` register.
- 0b1 This MSC supports internal priority partitioning and implements the internal priority (INTPRI) field in the `MPAMCFG_PRI` register.

If RIS is implemented, this field indicates that internal priority is implemented for the resource instance selected by `MPAMCFG_PART_SEL.RIS`.

Accessing the MPAMF_PRI_IDR:

This register is within the MPAM feature page memory frames. In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps.

`MPAMF_PRI_IDR` is read-only.

`MPAMF_PRI_IDR` must be readable from the Non-secure, Secure, Root, and Realm MPAM feature pages.

`MPAMF_PRI_IDR` is permitted to have the same contents when read from the Secure, Non-secure, Root, and Realm MPAM feature pages unless the register contents are different for the different versions:

- `MPAMF_PRI_IDR_s` is permitted to have either the same or different contents to `MPAMF_PRI_IDR_ns`, `MPAMF_PRI_IDR_rt`, or `MPAMF_PRI_IDR_rl`.
- `MPAMF_PRI_IDR_ns` is permitted to have either the same or different contents to `MPAMF_PRI_IDR_rt` or `MPAMF_PRI_IDR_rl`.

- MPAMF_PRI_IDR_rt is permitted to have either the same or different contents to MPAMF_PRI_IDR_rl.

There must be separate registers in the Secure (MPAMF_PRI_IDR_s), Non-secure (MPAMF_PRI_IDR_ns), Root (MPAMF_PRI_IDR_rt), and Realm (MPAMF_PRI_IDR_rl) MPAM feature pages.

When MPAMF_IDR.HAS_RIS is 1, MPAMF_PRI_IDR shows the configuration of priority partitioning for the resource instance selected by MPAMCFG_PART_SEL.RIS. Fields that mention RIS in their field descriptions have values that track the implemented properties of the resource instance. Fields that do not mention RIS are constant across all resource instances.

MPAMF_PRI_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0048	MPAMF_PRI_IDR_s

Accesses to this interface are RO.

MPAMF_PRI_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0048	MPAMF_PRI_IDR_ns

Accesses to this interface are RO.

MPAMF_PRI_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0048	MPAMF_PRI_IDR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MPAMF_PRI_IDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0048	MPAMF_PRI_IDR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.3.13 MPAMF_SIDR, MPAM Features Secure Identification Register

The MPAMF_SIDR characteristics are:

Purpose

The MPAMF_SIDR is a 32-bit read-only register that indicates the maximum Secure PARTID and Secure PMG on this MSC.

Configurations

The power domain of MPAMF_SIDR is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_SIDR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_SIDR is a 32-bit register.

Field descriptions



Bits [31:24]

Reserved, RES0.

S_PMG_MAX, bits [23:16]

Maximum value of Secure PMG supported by this component.

S_PARTID_MAX, bits [15:0]

Maximum value of Secure PARTID supported by this component.

Accessing the MPAMF_SIDR:

This register is only within the Secure MPAM feature page memory frame.

MPAMF_SIDR is read-only.

MPAMF_SIDR must only be readable from the Secure MPAM feature page. If the system or the MSC does not support the Secure address map, this register must not be accessible.

MPAMF_SIDR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0008	MPAMF_SIDR_s

Accesses to this interface are RO.

9.4 Memory-mapped partitioning configuration registers

This section lists the external partitioning configuration registers.

9.4.1 MPAMCFG_CASSOC, MPAM Cache Maximum Associativity Partition Configuration Register

The MPAMCFG_CASSOC characteristics are:

Purpose

The MPAMCFG_CASSOC is a 32-bit read/write register that controls the maximum fraction of the cache associativity that the PARTID selected by MPAMCFG_PART_SEL is permitted to allocate.

MPAMCFG_CASSOC_s controls the cache maximum associativity for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_CASSOC_ns controls the cache maximum associativity for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_CASSOC_rl controls the cache maximum associativity for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL. MPAMCFG_CASSOC_rt controls the cache maximum associativity for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_CASSOC is IMPLEMENTATION DEFINED.

This register is present only when MPAMF_IDR.HAS_CCAP_PART == 1, (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CCAP_IDR.HAS_CASSOC == 1. Otherwise, direct accesses to MPAMCFG_CASSOC are RES0.

Attributes

MPAMCFG_CASSOC is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

CASSOC, bits [15:0]

Maximum cache associativity usage in fixed-point fraction format by the partition selected by MPAMCFG_PART_SEL. The fraction represents the portion of the cache associativity that the PARTID is permitted to allocate. CASSOC controls the fraction of associativity in each associativity grouping of the cache. In a set associative cache, CASSOC applies to the fraction of the ways in each set.

The implemented width of the fixed-point fraction is given in MPAMF_CCAP_IDR.CASSOC_WD. Unimplemented bits within the field are RAZ/WI. The implemented bits of the CASSOC field are always the most significant bits of the field.

The fixed-point fraction CASSOC is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the cache that can be represented in an implementation with w implemented bits is 1.0 minus one half to the power w.

Accessing the MPAMCFG_CASSOC:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_CASSOC_s must only be accessible from the Secure MPAM feature page.

- MPAMCFG_CASSOC_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_CASSOC_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_CASSOC_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_CASSOC_s, MPAMCFG_CASSOC_ns, MPAMCFG_CASSOC_rt, and MPAMCFG_CASSOC_rl must be separate registers:

- The Secure instance (MPAMCFG_CASSOC_s) accesses the cache maximum associativity partitioning used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_CASSOC_ns) accesses the cache maximum associativity partitioning used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_CASSOC_rt) accesses the cache maximum associativity partitioning used for Root PARTIDs.
- The Realm instance (MPAMCFG_CASSOC_rl) accesses the cache maximum associativity partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_CASSOC access the cache maximum associativity partitioning configuration settings for the cache resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_CASSOC access the cache maximum associativity partitioning configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_CASSOC access the cache maximum associativity partitioning configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_CASSOC access the cache maximum associativity partitioning configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_CASSOC can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0118	MPAMCFG_CASSOC_s

Accesses to this interface are RW.

MPAMCFG_CASSOC can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0118	MPAMCFG_CASSOC_ns

Accesses to this interface are RW.

MPAMCFG_CASSOC can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0118	MPAMCFG_CASSOC_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_CASSOC can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0118	MPAMCFG_CASSOC_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.2 MPAMCFG_CMAX, MPAM Cache Maximum Capacity Partition Configuration Register

The MPAMCFG_CMAX characteristics are:

Purpose

The MPAMCFG_CMAX is a 32-bit read/write register that controls the maximum fraction of the cache capacity that the PARTID selected by MPAMCFG_PART_SEL is permitted to allocate.

MPAMCFG_CMAX_s controls the cache maximum capacity for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_CMAX_ns controls the cache maximum capacity for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_CMAX_rt controls the cache maximum capacity for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL. MPAMCFG_CMAX_rl controls the cache maximum capacity for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_CMAX is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_CCAP_PART == 1 and MPAMF_CCAP_IDR.NO_CMAX == 0. Otherwise, direct accesses to MPAMCFG_CMAX are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_CMAX is a 32-bit register.

Field descriptions



SOFTLIM, bit [31]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CCAP_IDR.HAS_CMAX_SOFTLIM == 1:

Soft limiting of CMAX. Soft limiting allows some allocations by a PARTID when its cache use is above the CMAX maximum cache capacity.

- 0b0 When CMAX cache capacity is exceeded, the partition is not allowed to increase its cache capacity usage. It is only permitted to replace a line that was previously occupied by a line allocated by that PARTID.
- 0b1 When CMAX cache capacity is exceeded, the partition is permitted to allocate capacity beyond CMAX, but only from invalid lines or lines belonging to disabled PARTIDs.

Otherwise:

Reserved, RES0.

Bits [30:16]

Reserved, RES0.

CMAX, bits [15:0]

Maximum cache capacity usage in fixed-point fraction format by the partition selected by MPAMCFG_PART_SEL. The fraction represents the portion of the total cache capacity that the PARTID is permitted to allocate.

The implemented width of the fixed-point fraction is given in `MPAMF_CCAP_IDR.CMAX_WD`. Unimplemented bits within the field are RAZ/WI. The implemented bits of the CMAX field are always the most significant bits of the field.

The fixed-point fraction CMAX is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the cache that can be represented in an implementation with `w` implemented bits is 1.0 minus one half to the power `w`.

Accessing the MPAMCFG_CMAX:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MPAMCFG_CMAX_s` must only be accessible from the Secure MPAM feature page.
- `MPAMCFG_CMAX_ns` must only be accessible from the Non-secure MPAM feature page.
- `MPAMCFG_CMAX_rt` must only be accessible from the Root MPAM feature page.
- `MPAMCFG_CMAX_rl` must only be accessible from the Realm MPAM feature page.

`MPAMCFG_CMAX_s`, `MPAMCFG_CMAX_ns`, `MPAMCFG_CMAX_rt`, and `MPAMCFG_CMAX_rl` must be separate registers:

- The Secure instance (`MPAMCFG_CMAX_s`) accesses the cache capacity partitioning used for Secure PARTIDs.
- The Non-secure instance (`MPAMCFG_CMAX_ns`) accesses the cache capacity partitioning used for Non-secure PARTIDs.
- The Root instance (`MPAMCFG_CMAX_rt`) accesses the cache capacity partitioning used for Root PARTIDs.
- The Realm instance (`MPAMCFG_CMAX_rl`) accesses the cache capacity partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to `MPAMCFG_CMAX` access the cache maximum capacity partitioning configuration settings for the cache resource instance selected by `MPAMCFG_PART_SEL.RIS` and the PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`.

When RIS is not implemented, loads and stores to `MPAMCFG_CMAX` access the cache maximum capacity partitioning configuration settings for the PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`.

When PARTID narrowing is implemented, loads and stores to `MPAMCFG_CMAX` access the cache maximum capacity partitioning configuration settings for the internal PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`, and `MPAMCFG_PART_SEL.INTERNAL` must be 1.

When PARTID narrowing is not implemented, loads and stores to `MPAMCFG_CMAX` access the cache maximum capacity partitioning configuration settings for the request PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`, and `MPAMCFG_PART_SEL.INTERNAL` must be 0.

`MPAMCFG_CMAX` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_s</code>	0x0108	<code>MPAMCFG_CMAX_s</code>

Accesses to this interface are RW.

MPAMCFG_CMAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0108	MPAMCFG_CMAX_ns

Accesses to this interface are RW.

MPAMCFG_CMAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0108	MPAMCFG_CMAX_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_CMAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0108	MPAMCFG_CMAX_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.3 MPAMCFG_CMIN, MPAM Cache Minimum Capacity Partition Configuration Register

The MPAMCFG_CMIN characteristics are:

Purpose

The MPAMCFG_CMIN is a 32-bit read/write register that controls the fraction of the cache capacity that the PARTID selected by MPAMCFG_PART_SEL has priority to allocate.

MPAMCFG_CMIN_s controls the cache minimum capacity for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_CMIN_ns controls the cache minimum capacity for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_CMIN_rl controls the cache minimum capacity for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL. MPAMCFG_CMIN_rt controls the cache minimum capacity for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_CMIN is IMPLEMENTATION DEFINED.

This register is present only when MPAMF_IDR.HAS_CCAP_PART == 1, (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CCAP_IDR.HAS_CMIN == 1. Otherwise, direct accesses to MPAMCFG_CMIN are RES0.

Attributes

MPAMCFG_CMIN is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

CMIN, bits [15:0]

Minimum cache capacity usage in fixed-point fraction format by the partition selected by MPAMCFG_PART_SEL. The fraction represents the portion of the total cache capacity that the PARTID has priority to allocate.

The implemented width of the fixed-point fraction is the same as the width of MPAMCFG_CMAX.CMAX which is given in MPAMF_CCAP_IDR.CMAX_WD. Unimplemented bits within the field are RAZ/WI. The implemented bits of the CMIN field are always the most significant bits of the field.

The fixed-point fraction CMIN is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the cache that can be represented in an implementation with w implemented bits is 1.0 minus one half to the power w.

Accessing the MPAMCFG_CMIN:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_CMIN_s must only be accessible from the Secure MPAM feature page.

- MPAMCFG_CMIN_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_CMIN_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_CMIN_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_CMIN_s, MPAMCFG_CMIN_ns, MPAMCFG_CMIN_rt, and MPAMCFG_CMIN_rl must be separate registers:

- The Secure instance (MPAMCFG_CMIN_s) accesses the cache minimum capacity partitioning used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_CMIN_ns) accesses the cache minimum capacity partitioning used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_CMIN_rt) accesses the cache minimum capacity partitioning used for Root PARTIDs.
- The Realm instance (MPAMCFG_CMIN_rl) accesses the cache minimum capacity partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_CMIN access the cache minimum capacity partitioning configuration settings for the cache resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_CMIN access the cache minimum capacity partitioning configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_CMIN access the cache minimum capacity partitioning configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_CMIN access the cache minimum capacity partitioning configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_CMIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0110	MPAMCFG_CMIN_s

Accesses to this interface are RW.

MPAMCFG_CMIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0110	MPAMCFG_CMIN_ns

Accesses to this interface are RW.

MPAMCFG_CMIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0110	MPAMCFG_CMIN_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_CMIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0110	MPAMCFG_CMIN_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.4 MPAMCFG_CPBM<n>, MPAM Cache Portion Bitmap Partition Configuration Register, n = 0 - 1023

The MPAMCFG_CPBM<n> characteristics are:

Purpose

The MPAMCFG_CPBM<n> register array gives access to the cache portion bitmap. Each register in the array is a read/write register that configures the cache portions numbered from <n * 32> to <31 + (n * 32)> that a PARTID is allowed to allocate.

After setting [MPAMCFG_PART_SEL](#) with a PARTID, software writes to the MPAMCFG_CPBM<n> register to configure which cache portions the PARTID is allowed to allocate.

The MPAMCFG_CPBM<n> register that contains the bitmap bit corresponding to cache portion p has n equal to $\lfloor p/32 \rfloor$. The field, P<x>, of that MPAMCFG_CPBM<n> register that contains the bitmap bit corresponding to cache portion p has x equal to $p \bmod 32$.

MPAMCFG_CPBM<n>_s controls cache portions for the Secure PARTID selected by the Secure instance of [MPAMCFG_PART_SEL](#). MPAMCFG_CPBM<n>_ns controls the cache portions for the Non-secure PARTID selected by the Non-secure instance of [MPAMCFG_PART_SEL](#). MPAMCFG_CPBM<n>_rt controls cache portions for the Root PARTID selected by the Root instance of [MPAMCFG_PART_SEL](#). MPAMCFG_CPBM<n>_rl controls the cache portions for the Realm PARTID selected by the Non-secure instance of [MPAMCFG_PART_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the control settings accessed are those of the resource instance currently selected by [MPAMCFG_PART_SEL.RIS](#) and the PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#).

Configurations

The power domain of MPAMCFG_CPBM<n> is IMPLEMENTATION DEFINED.

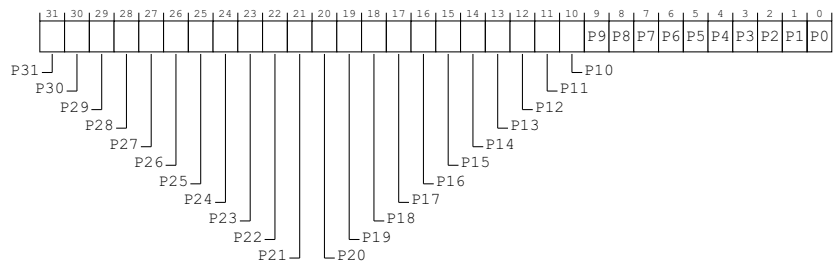
This register is present only when [FEAT_MPAM](#) is implemented and [MPAMF_IDR.HAS_CPOR_PART](#) == 1. Otherwise, direct accesses to MPAMCFG_CPBM<n> are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_CPBM<n> is a 32-bit register.

Field descriptions



P<x>, bit [x], for x = 31 to 0

Portion allocation control bit. Each cache portion allocation control bit, MPAMCFG_CPBM<n>.P<x>, grants permission to the PARTID selected by [MPAMCFG_PART_SEL](#) to allocate cache lines within cache portion <n*32> + x.

- 0b0 The PARTID is not permitted to allocate into cache portion <n * 32> + x.
- 0b1 The PARTID is permitted to allocate within cache portion <n * 32> + x.

The number of bits in the cache portion partitioning bit map of this component is given in [MPAMF_CPOR_IDR.CPBM_WD](#). [MPAMF_CPOR_IDR.CPBM_WD](#) contains a value from 1 to 2¹⁵, inclusive. Values of [MPAMF_CPOR_IDR.CPBM_WD](#) greater than 32 require an array of 32-bit [MPAMCFG_CPBM<n>](#) registers to access the cache portion bitmap, up to 1024 registers.

When $(n * 32) + x > \text{UInt}(\text{MPAMF_CPOR_IDR.CPBM_WD})$, access to this field is RES0.

Accessing the MPAMCFG_CPBM<n>:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- [MPAMCFG_CPBM<n>_s](#) must only be accessible from the Secure MPAM feature page.
- [MPAMCFG_CPBM<n>_ns](#) must only be accessible from the Non-secure MPAM feature page.
- [MPAMCFG_CPBM<n>_rt](#) must only be accessible from the Root MPAM feature page.
- [MPAMCFG_CPBM<n>_rl](#) must only be accessible from the Realm MPAM feature page.

[MPAMCFG_CPBM<n>_s](#), [MPAMCFG_CPBM<n>_ns](#), [MPAMCFG_CPBM<n>_rt](#), and [MPAMCFG_CPBM<n>_rl](#) must be separate registers:

- The Secure instance ([MPAMCFG_CPBM<n>_s](#)) accesses the cache portion bitmap used for Secure PARTIDs.
- The Non-secure instance ([MPAMCFG_CPBM<n>_ns](#)) accesses the cache portion bitmap used for Non-secure PARTIDs.
- The Root instance ([MPAMCFG_CPBM<n>_rt](#)) accesses the cache portion bitmap used for Root PARTIDs.
- The Realm instance ([MPAMCFG_CPBM<n>_rl](#)) accesses the cache portion bitmap used for Realm PARTIDs.

When RIS is implemented, loads and stores to [MPAMCFG_CPBM<n>](#) access the cache portion bitmap configuration settings for the cache resource instance selected by [MPAMCFG_PART_SEL.RIS](#) and the PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#).

When RIS is not implemented, loads and stores to [MPAMCFG_CPBM<n>](#) access the cache portion bitmap configuration settings for the PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#).

When PARTID narrowing is implemented, loads and stores to [MPAMCFG_CPBM<n>](#) access the cache portion bitmap configuration settings for the internal PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#), and [MPAMCFG_PART_SEL.INTERNAL](#) must be 1.

When PARTID narrowing is not implemented, loads and stores to [MPAMCFG_CPBM<n>](#) access the cache portion bitmap configuration settings for the request PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#), and [MPAMCFG_PART_SEL.INTERNAL](#) must be 0.

[MPAMCFG_CPBM<n>](#) can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	$0x1000 + (4 * n)$	MPAMCFG_CPBM<n>_s

Accesses to this interface are RW.

MPAMCFG_CPBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x1000 + (4 * n)	MPAMCFG_CPBM<n>_ns

Accesses to this interface are RW.

MPAMCFG_CPBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x1000 + (4 * n)	MPAMCFG_CPBM<n>_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_CPBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x1000 + (4 * n)	MPAMCFG_CPBM<n>_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.5 MPAMCFG_DIS, MPAM Partition Configuration Disable Register

The MPAMCFG_DIS characteristics are:

Purpose

Disables a PARTID configuration as set in other MPAMCFG registers.

MPAMCFG_DIS_s disables a Secure PARTID. MPAMCFG_DIS_ns disables a Non-secure PARTID. MPAMCFG_DIS_rl disables a Realm PARTID. MPAMCFG_DIS_rt disables a Root PARTID.

Configurations


The power domain of MPAMCFG_DIS is IMPLEMENTATION DEFINED.

This register is present only when (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_IDR.HAS_ENDIS == 1. Otherwise, direct accesses to MPAMCFG_DIS are RES0.

Attributes

MPAMCFG_DIS is a 32-bit register.

Field descriptions



NFU, bit [31]

When MPAMF_IDR.HAS_NFU == 1:

No Future Use.

0b0 Control settings of the disabled PARTID must be retained.

0b1 Control settings of the disabled PARTID may take an UNKNOWN value.

Otherwise:

Reserved, RES0.

Bits [30:16]

Reserved, RES0.

PARTID, bits [15:0]

Selects the PARTID to disable.

Accessing the MPAMCFG_DIS:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_DIS_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_DIS_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_DIS_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_DIS_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_DIS_s, MPAMCFG_DIS_ns, MPAMCFG_DIS_rt, and MPAMCFG_DIS_rl must be separate registers:

- The Secure instance (MPAMCFG_DIS_s) accesses the PARTID disable used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_DIS_ns) accesses the PARTID disable used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_DIS_rt) accesses the PARTID disable used for Root PARTIDs.
- The Realm instance (MPAMCFG_DIS_rl) accesses the PARTID disable used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_DIS access the PARTID disable configuration settings for the PARTID disable resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_DIS access the PARTID disable configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_DIS access the PARTID disable configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_DIS access the PARTID disable configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_DIS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0310	MPAMCFG_DIS_s

Accesses to this interface are WO/RAZ.

MPAMCFG_DIS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0310	MPAMCFG_DIS_ns

Accesses to this interface are WO/RAZ.

MPAMCFG_DIS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0310	MPAMCFG_DIS_rt

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

MPAMCFG_DIS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0310	MPAMCFG_DIS_rl

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

9.4.6 MPAMCFG_EN, MPAM Partition Configuration Enable Register

The MPAMCFG_EN characteristics are:

Purpose

Enables a PARTID configuration as set in other MPAMCFG registers.

MPAMCFG_EN_s enables a Secure PARTID. MPAMCFG_EN_ns enables a Non-secure PARTID. MPAMCFG_EN_rl enables a Realm PARTID. MPAMCFG_EN_rt enables a Root PARTID.

Configurations


The power domain of MPAMCFG_EN is IMPLEMENTATION DEFINED.

This register is present only when (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_IDR.HAS_ENDIS == 1. Otherwise, direct accesses to MPAMCFG_EN are RES0.

Attributes

MPAMCFG_EN is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

PARTID, bits [15:0]

Selects the PARTID to enable.

Accessing the MPAMCFG_EN:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_EN_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_EN_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_EN_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_EN_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_EN_s, MPAMCFG_EN_ns, MPAMCFG_EN_rt, and MPAMCFG_EN_rl must be separate registers:

- The Secure instance (MPAMCFG_EN_s) accesses the PARTID enable used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_EN_ns) accesses the PARTID enable used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_EN_rt) accesses the PARTID enable used for Root PARTIDs.
- The Realm instance (MPAMCFG_EN_rl) accesses the PARTID enable used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_EN access the PARTID enable configuration settings for the PARTID enable resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_EN access the PARTID enable configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_EN access the PARTID enable configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_EN access the PARTID enable configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_EN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0300	MPAMCFG_EN_s

Accesses to this interface are WO/RAZ.

MPAMCFG_EN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0300	MPAMCFG_EN_ns

Accesses to this interface are WO/RAZ.

MPAMCFG_EN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0300	MPAMCFG_EN_rt

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

MPAMCFG_EN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0300	MPAMCFG_EN_rl

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

9.4.7 MPAMCFG_EN_FLAGS, MPAM Partition Configuration Enable Flags Register

The MPAMCFG_EN_FLAGS characteristics are:

Purpose

Enable flags for 32 PARTIDs.

MPAMCFG_EN_FLAGS_s gives read/write access to 32 Secure PARTIDs.

MPAMCFG_EN_FLAGS_ns gives read/write access to 32 Non-secure PARTIDs.

MPAMCFG_EN_FLAGS_rl gives read/write access to 32 Realm PARTIDs.

MPAMCFG_EN_FLAGS_rt gives read/write access to 32 Root PARTIDs.

Configurations

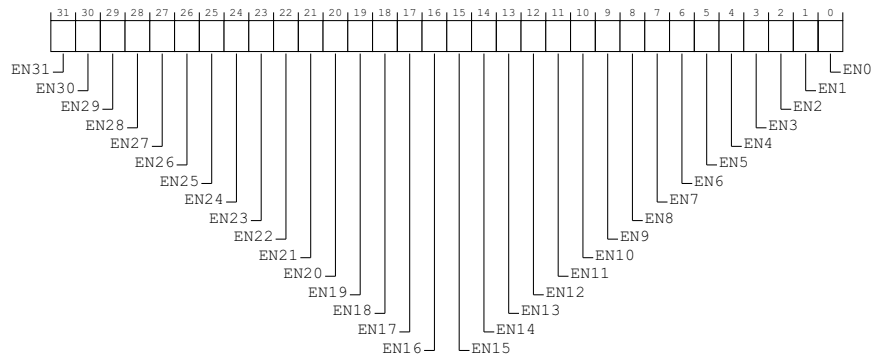
The power domain of MPAMCFG_EN_FLAGS is IMPLEMENTATION DEFINED.

This register is present only when (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_IDR.HAS_ENDIS == 1. Otherwise, direct accesses to MPAMCFG_EN_FLAGS are RES0.

Attributes

MPAMCFG_EN_FLAGS is a 32-bit register.

Field descriptions



EN<x>, bit [x], for x = 31 to 0

PARTID Enable flags. The group of flags accessed is selected by MPAMCFG_PART_SEL.PARTID_SEL & 0xFFE0 in bit [0] to (MPAMCFG_PART_SEL.PARTID_SEL & 0xFFE0) + 31 in bit [31].

0b0 The PARTID is disabled.

0b1 The PARTID is enabled.

Each bit in MPAMCFG_EN_FLAGS gives access to the same state as controlled by MPAMCFG_EN and MPAMCFG_DIS.

Bits MPAMCFG_EN_FLAGS.EN<x>, where (MPAMCFG_PART_SEL.PARTID_SEL & 0xFFE0) + x is greater than MPAMF_IDR.PARTID_MAX, are not required to be implemented.

As with other partitioning controls, the enable flag for PARTID 0 must be reset to 0b1 (enabled).

Accessing the MPAMCFG_EN_FLAGS:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_EN_FLAGS_s must only be accessible from the Secure MPAM feature page.

- MPAMCFG_EN_FLAGS_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_EN_FLAGS_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_EN_FLAGS_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_EN_FLAGS_s, MPAMCFG_EN_FLAGS_ns, MPAMCFG_EN_FLAGS_rt, and MPAMCFG_EN_FLAGS_rl must be separate registers:

- The Secure instance (MPAMCFG_EN_FLAGS_s) accesses the PARTID enable used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_EN_FLAGS_ns) accesses the PARTID enable used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_EN_FLAGS_rt) accesses the PARTID enable used for Root PARTIDs.
- The Realm instance (MPAMCFG_EN_FLAGS_rl) accesses the PARTID enable used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_EN_FLAGS access the PARTID enable configuration settings for the PARTID enable resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_EN_FLAGS access the PARTID enable configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_EN_FLAGS access the PARTID enable configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_EN_FLAGS access the PARTID enable configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_EN_FLAGS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0320	MPAMCFG_EN_FLAGS_s

Accesses to this interface are RW.

MPAMCFG_EN_FLAGS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0320	MPAMCFG_EN_FLAGS_ns

Accesses to this interface are RW.

MPAMCFG_EN_FLAGS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0320	MPAMCFG_EN_FLAGS_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_EN_FLAGS can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0320	MPAMCFG_EN_FLAGS_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.8 MPAMCFG_INTPARTID, MPAM Internal PARTID Narrowing Configuration Register

The MPAMCFG_INTPARTID characteristics are:

Purpose

MPAMCFG_INTPARTID is a 32-bit read/write register that controls the mapping of the PARTID selected by MPAMCFG_PART_SEL into a narrower internal PARTID (intPARTID).

MPAMCFG_INTPARTID_s controls the mapping for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_INTPARTID_ns controls the mapping for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL.

MPAMCFG_INTPARTID_rt controls the mapping for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL. MPAMCFG_INTPARTID_rl controls the mapping for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

The MPAMCFG_INTPARTID register associates the request PARTID (reqPARTID) in the MPAMCFG_PART_SEL register with an internal PARTID (intPARTID) in this register. To set that association, store reqPARTID into the MPAMCFG_PART_SEL register and then store the intPARTID into the MPAMCFG_INTPARTID register. To read the association, store reqPARTID into the MPAMCFG_PART_SEL register and then read MPAMCFG_INTPARTID.

If the intPARTID stored into MPAMCFG_INTPARTID is out-of-range or does not have the INTERNAL bit set, the association of reqPARTID to intPARTID is not written and MPAMF_ESR is set to indicate an intPARTID_Range error.

If MPAMCFG_PART_SEL.INTERNAL is 1 when MPAMCFG_INTPARTID is read or written, MPAMF_ESR is set to indicate an Unexpected_INTERNAL error.

Configurations

The power domain of MPAMCFG_INTPARTID is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_PARTID_NRW == 1. Otherwise, direct accesses to MPAMCFG_INTPARTID are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_INTPARTID is a 32-bit register.

Field descriptions



Bits [31:17]

Reserved, RES0.

INTERNAL, bit [16]

Internal PARTID flag.

This bit must be 1 when written to the register. If written as 0, the write will not update the reqPARTID to intPARTID association.

On a read of this register, the bit will always read the value last written.

INTPARTID, bits [15:0]

This field contains the intPARTID mapped to the reqPARTID in MPAMCFG_PART_SEL.

The maximum intPARTID supported is MPAMF_PARTID_NRW_IDR.INTPARTID_MAX.

Accessing the MPAMCFG_INTPARTID:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_INTPARTID_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_INTPARTID_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_INTPARTID_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_INTPARTID_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_INTPARTID_s, MPAMCFG_INTPARTID_ns, MPAMCFG_INTPARTID_rt, and MPAMCFG_INTPARTID_rl must be separate registers:

- The Secure instance (MPAMCFG_INTPARTID_s) accesses the PARTID narrowing used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_INTPARTID_ns) accesses the PARTID narrowing used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_INTPARTID_rt) accesses the PARTID narrowing used for Root PARTIDs.
- The Realm instance (MPAMCFG_INTPARTID_rl) accesses the PARTID narrowing used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_INTPARTID access the PARTID narrowing configuration settings without being affected by MPAMCFG_PART_SEL.RIS.

Loads and stores to MPAMCFG_INTPARTID access the PARTID narrowing configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_INTPARTID can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0600	MPAMCFG_INTPARTID_s

Accesses to this interface are RW.

MPAMCFG_INTPARTID can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0600	MPAMCFG_INTPARTID_ns

Accesses to this interface are RW.

MPAMCFG_INTPARTID can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0600	MPAMCFG_INTPARTID_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_INTPARTID can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0600	MPAMCFG_INTPARTID_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.9 MPAMCFG_MBW_MAX, MPAM Memory Bandwidth Maximum Partition Configuration Register

The MPAMCFG_MBW_MAX characteristics are:

Purpose

MPAMCFG_MBW_MAX is a 32-bit read/write register that controls the maximum fraction of memory bandwidth that the PARTID selected by MPAMCFG_PART_SEL is permitted to use.

MPAMCFG_MBW_MAX_s controls maximum bandwidth for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_MBW_MAX_ns controls the maximum bandwidth for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_MBW_MAX_rt controls the maximum bandwidth for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_MAX_rl controls the maximum bandwidth for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

A PARTID that has used more than MAX is given no access to additional bandwidth if HARDLIM == 1 or is given additional bandwidth only if there are no requests from PARTIDs that have not exceeded their MAX if HARDLIM == 0.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_MBW_MAX is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MBW_PART == 1 and MPAMF_MBW_IDR.HAS_MAX == 1. Otherwise, direct accesses to MPAMCFG_MBW_MAX are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_MBW_MAX is a 32-bit register.

Field descriptions



HARDLIM, bit [31]

Maximum-bandwidth limit behavior selection.

- 0b0 Soft limit: when MAX bandwidth is exceeded, the partition contends with a low preference for downstream bandwidth beyond MAX.
- 0b1 Hard limit: when MAX bandwidth is exceeded, the partition does not use any more bandwidth until the memory bandwidth measurement for the partition falls below MAX.

Accessing this field has the following behavior:

- When MPAMF_MBW_IDR.MAX_LIM == 0b00, access to this field is RW.
- When MPAMF_MBW_IDR.MAX_LIM == 0b01, access to this field is RAZ/WI.
- When MPAMF_MBW_IDR.MAX_LIM == 0b10, access to this field is RAO/WI.

Bits [30:16]

Reserved, RES0.

MAX, bits [15:0]

Memory maximum bandwidth allocated to the partition selected by `MPAMCFG_PART_SEL`. MAX is in fixed-point fraction format. The fraction represents the portion of the total memory bandwidth capacity through the controlled component that the PARTID is permitted to allocate.

The implemented width of the fixed-point fraction is given in `MPAMF_MBW_IDR.BWA_WD`. Unimplemented bits are RAZ/WI. The implemented bits of the MAX field are always to the left of the field. For example, if `BWA_WD = 3`, the implemented bits are `MPAMCFG_MBW_MAX[15:13]` and `MPAMCFG_MBW_MAX[12:0]` are unimplemented.

The fixed-point fraction MAX is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the bandwidth that can be represented in an implementation with `w` implemented bits is 1.0 minus one half to the power `w`.

Accessing the MPAMCFG_MBW_MAX:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MPAMCFG_MBW_MAX_s` must only be accessible from the Secure MPAM feature page.
- `MPAMCFG_MBW_MAX_ns` must only be accessible from the Non-secure MPAM feature page.
- `MPAMCFG_MBW_MAX_rt` must only be accessible from the Root MPAM feature page.
- `MPAMCFG_MBW_MAX_rl` must only be accessible from the Realm MPAM feature page.

`MPAMCFG_MBW_MAX_s`, `MPAMCFG_MBW_MAX_ns`, `MPAMCFG_MBW_MAX_rt`, and `MPAMCFG_MBW_MAX_rl` must be separate registers:

- The Secure instance (`MPAMCFG_MBW_MAX_s`) accesses the memory maximum bandwidth partitioning used for Secure PARTIDs.
- The Non-secure instance (`MPAMCFG_MBW_MAX_ns`) accesses the memory maximum bandwidth partitioning used for Non-secure PARTIDs.
- The Root instance (`MPAMCFG_MBW_MAX_rt`) accesses the memory maximum bandwidth partitioning used for Root PARTIDs.
- The Realm instance (`MPAMCFG_MBW_MAX_rl`) accesses the memory maximum bandwidth partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to `MPAMCFG_MBW_MAX` access the memory maximum bandwidth partitioning configuration settings for the bandwidth resource instance selected by `MPAMCFG_PART_SEL.RIS` and the PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`.

When RIS is not implemented, loads and stores to `MPAMCFG_MBW_MAX` access the memory maximum bandwidth partitioning configuration settings for the PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`.

When PARTID narrowing is implemented, loads and stores to `MPAMCFG_MBW_MAX` access the memory maximum bandwidth partitioning configuration settings for the internal PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`, and `MPAMCFG_PART_SEL.INTERNAL` must be 1.

When PARTID narrowing is not implemented, loads and stores to `MPAMCFG_MBW_MAX` access the memory maximum bandwidth partitioning configuration settings for the request PARTID selected by `MPAMCFG_PART_SEL.PARTID_SEL`, and `MPAMCFG_PART_SEL.INTERNAL` must be 0.

`MPAMCFG_MBW_MAX` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	<code>MPAMF_BASE_s</code>	<code>0x0208</code>	<code>MPAMCFG_MBW_MAX_s</code>

Accesses to this interface are RW.

MPAMCFG_MBW_MAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0208	MPAMCFG_MBW_MAX_ns

Accesses to this interface are RW.

MPAMCFG_MBW_MAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0208	MPAMCFG_MBW_MAX_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_MBW_MAX can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0208	MPAMCFG_MBW_MAX_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.10 MPAMCFG_MBW_MIN, MPAM Memory Bandwidth Minimum Partition Configuration Register

The MPAMCFG_MBW_MIN characteristics are:

Purpose

MPAMCFG_MBW_MIN is a 32-bit read/write register that controls the minimum fraction of memory bandwidth that the PARTID selected by MPAMCFG_PART_SEL is permitted to use.

MPAMCFG_MBW_MIN_s controls the minimum bandwidth for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_MBW_MIN_ns controls the minimum bandwidth for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_MBW_MIN_rt controls the minimum bandwidth for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_MIN_rl controls the minimum bandwidth for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

A PARTID that has used less than MIN is given preferential access to bandwidth.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_MBW_MIN is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MBW_PART == 1 and MPAMF_MBW_IDR.HAS_MIN == 1. Otherwise, direct accesses to MPAMCFG_MBW_MIN are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_MBW_MIN is a 32-bit register.

Field descriptions



Bits [31:16]

Reserved, RES0.

MIN, bits [15:0]

Memory minimum bandwidth allocated to the partition selected by MPAMCFG_PART_SEL. MIN is in fixed-point fraction format. The fraction represents the portion of the total memory bandwidth capacity through the controlled component that the PARTID is permitted to allocate.

The implemented width of the fixed-point fraction is given in MPAMF_MBW_IDR.BWA_WD. Unimplemented bits are RAZ/WI. The implemented bits of the MIN field are always to the left of the field. For example, if BWA_WD = 4, the implemented bits are MPAMCFG_MBW_MIN[15:12] and MPAMCFG_MBW_MIN[11:0] are unimplemented.

The fixed-point fraction MIN is less than 1. The implied binary point is between bits 15 and 16. This representation has as the largest fraction of the bandwidth that can be represented in an implementation with w implemented bits is 1.0 minus one half to the power w.

Accessing the MPAMCFG_MBW_MIN:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_MBW_MIN_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_MBW_MIN_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_MBW_MIN_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_MBW_MIN_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_MBW_MIN_s, MPAMCFG_MBW_MIN_ns, MPAMCFG_MBW_MIN_rt, and MPAMCFG_MBW_MIN_rl must be separate registers:

- The Secure instance (MPAMCFG_MBW_MIN_s) accesses the memory minimum bandwidth partitioning used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_MBW_MIN_ns) accesses the memory minimum bandwidth partitioning used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_MBW_MIN_rt) accesses the memory minimum bandwidth partitioning used for Root PARTIDs.
- The Realm instance (MPAMCFG_MBW_MIN_rl) accesses the memory minimum bandwidth partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_MBW_MIN access the memory minimum bandwidth partitioning configuration settings for the bandwidth resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_MBW_MIN access the memory minimum bandwidth partitioning configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_MBW_MIN access the memory minimum bandwidth partitioning configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_MBW_MIN access the memory minimum bandwidth partitioning configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_MBW_MIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0200	MPAMCFG_MBW_MIN_s

Accesses to this interface are RW.

MPAMCFG_MBW_MIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0200	MPAMCFG_MBW_MIN_ns

Accesses to this interface are RW.

MPAMCFG_MBW_MIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0200	MPAMCFG_MBW_MIN_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_MBW_MIN can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0200	MPAMCFG_MBW_MIN_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.11 MPAMCFG_MBW_PBM<n>, MPAM Bandwidth Portion Bitmap Partition Configuration Register, n = 0 - 127

The MPAMCFG_MBW_PBM<n> characteristics are:

Purpose

The MPAMCFG_MBW_PBM<n> register array gives access to the memory bandwidth portion bitmap. Each register in the array is a read/write register that configures whether a PARTID is allowed to allocate bandwidth portions within a range.

The range of portions covered in MPAMCFG_MBW_PBM<n> is from portion <32*n> to portion <32*n +31>.

After setting [MPAMCFG_PART_SEL](#) with a PARTID, software writes to one or more of the MPAMCFG_MBW_PBM<n> registers to configure with bandwidth portions the PARTID is allowed to allocate.

The MPAMCFG_MBW_PBM<n> register that contains the bitmap bit corresponding to memory bandwidth portion p has n equal to p[11:5]. The field, P<x> of that MPAMCFG_MBW_PBM<n> register that contains the bitmap bit corresponding to memory bandwidth portion p has <x> equal to p[4:0].

The MPAMCFG_MBW_PBM<n>_s registers control the bandwidth portion bitmap for the Secure PARTID selected by the Secure instance of [MPAMCFG_PART_SEL](#). The MPAMCFG_MBW_PBM<n>_ns registers control the bandwidth portion bitmap for the Non-secure PARTID selected by the Non-secure instance of [MPAMCFG_PART_SEL](#). The MPAMCFG_MBW_PBM<n>_rt registers control the bandwidth portion bitmap for the Root PARTID selected by the Root instance of [MPAMCFG_PART_SEL](#). The MPAMCFG_MBW_PBM<n>_rl registers control the bandwidth portion bitmap for the Realm PARTID selected by the Realm instance of [MPAMCFG_PART_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the control settings accessed are those of the resource instance currently selected by [MPAMCFG_PART_SEL.RIS](#) and the PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#).

Configurations

The power domain of MPAMCFG_MBW_PBM<n> is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MBW_PART](#) == 1 and [MPAMF_MBW_IDR.HAS_PBM](#) == 1. Otherwise, direct accesses to MPAMCFG_MBW_PBM<n> are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_MBW_PBM<n> is a 32-bit register.

Field descriptions



P<x>, bit [x], for x = 31 to 0

Portion allocation control bit. Each bandwidth portion allocation control bit MPAMCFG_MBW_PBM<n>.P<x> grants permission to the PARTID selected by MPAMCFG_PART_SEL to allocate bandwidth within bandwidth portion <32*n> + <x>.

0b0 The PARTID is not permitted to allocate into bandwidth portion <32*n> + <x>.

0b1 The PARTID is permitted to allocate within bandwidth portion <32*n> + <x>.

The number of bits in the bandwidth portion partitioning bit map of this component is given in MPAMF_MBW_IDR.BWPBM_WD. BWPBM_WD contains a value from 1 to 2¹², inclusive. Values of MPAMF_MBW_IDR.BWPBM_WD greater than 32 require a group of 32-bit registers to access the bandwidth portion bitmap, up to 128 32-bit registers.

When $(n * 32) + x > \text{UInt}(\text{MPAMF_MBW_IDR.BWPBM_WD})$, access to this field is RES0.

Accessing the MPAMCFG_MBW_PBM<n>:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_MBW_PBM<n>_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_MBW_PBM<n>_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_MBW_PBM<n>_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_MBW_PBM<n>_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_MBW_PBM<n>_s, MPAMCFG_MBW_PBM<n>_ns, MPAMCFG_MBW_PBM<n>_rt, and MPAMCFG_MBW_PBM<n>_rl must be separate registers:

- The Secure instance (MPAMCFG_MBW_PBM<n>_s) accesses the memory bandwidth portion bitmap used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_MBW_PBM<n>_ns) accesses the memory bandwidth portion bitmap used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_MBW_PBM<n>_rt) accesses the memory bandwidth portion bitmap used for Root PARTIDs.
- The Realm instance (MPAMCFG_MBW_PBM<n>_rl) accesses the memory bandwidth portion bitmap used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_MBW_PBM<n> access the memory bandwidth portion bitmap configuration settings for the bandwidth resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_MBW_PBM<n> access the memory bandwidth portion bitmap configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_MBW_PBM<n> access the memory bandwidth portion bitmap configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_MBW_PBM<n> access the memory bandwidth portion bitmap configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_MBW_PBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x2000 + (4 * n)	MPAMCFG_MBW_PBM<n>_s

Accesses to this interface are RW.

MPAMCFG_MBW_PBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x2000 + (4 * n)	MPAMCFG_MBW_PBM<n>_ns

Accesses to this interface are RW.

MPAMCFG_MBW_PBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x2000 + (4 * n)	MPAMCFG_MBW_PBM<n>_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_MBW_PBM<n> can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x2000 + (4 * n)	MPAMCFG_MBW_PBM<n>_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.12 MPAMCFG_MBW_PROP, MPAM Memory Bandwidth Proportional Stride Partition Configuration Register

The MPAMCFG_MBW_PROP characteristics are:

Purpose

Controls the proportional stride of memory bandwidth that the PARTID selected by [MPAMCFG_PART_SEL](#) uses.

MPAMCFG_MBW_PROP_s controls the bandwidth proportional stride for the Secure PARTID selected by the Secure instance of [MPAMCFG_PART_SEL](#). MPAMCFG_MBW_PROP_ns controls the bandwidth proportional stride for the Non-secure PARTID selected by the Non-secure instance of [MPAMCFG_PART_SEL](#). MPAMCFG_MBW_PROP_rt controls the bandwidth proportional stride for the Root PARTID selected by the Root instance of [MPAMCFG_PART_SEL](#). MPAMCFG_MBW_PROP_rl controls the bandwidth proportional stride for the Realm PARTID selected by the Realm instance of [MPAMCFG_PART_SEL](#).

Proportional stride is a relative cost of bandwidth requested by one PARTID in relation to the costs of the bandwidths requested by each other PARTID also competing to use the bandwidth.

If [MPAMF_IDR.HAS_RIS](#) is 1, the control settings accessed are those of the resource instance currently selected by [MPAMCFG_PART_SEL.RIS](#) and the PARTID selected by [MPAMCFG_PART_SEL.PARTID_SEL](#).

Configurations

The power domain of MPAMCFG_MBW_PROP is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MBW_PART](#) == 1 and [MPAMF_MBW_IDR.HAS_PROP](#) == 1. Otherwise, direct accesses to MPAMCFG_MBW_PROP are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_MBW_PROP is a 32-bit register.

Field descriptions



EN, bit [31]

Enable proportional stride bandwidth partitioning.

- 0b0 The selected partition is not regulated by proportional stride bandwidth partitioning.
- 0b1 The selected partition has bandwidth usage regulated by proportional stride bandwidth partitioning as controlled by STRIDEM1.

Bits [30:16]

Reserved, RES0.

STRIDEM1, bits [15:0]

Memory bandwidth stride minus 1 allocated to the partition selected by [MPAMCFG_PART_SEL](#). STRIDEM1 represents the normalized cost of bandwidth consumption by the partition.

The proportional stride partitioning control parameter is an unsigned integer representing the normalized cost to a partition for consuming bandwidth. Larger values have a larger cost and correspond to a lesser allocation of bandwidth while smaller values indicate a lesser cost and therefore a higher allocation of bandwidth.

The implemented width of STRIDEM1 is given in [MPAMF_MBW_IDR.BWA_WD](#).

Accessing the MPAMCFG_MBW_PROP:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_MBW_PROP_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_MBW_PROP_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_MBW_PROP_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_MBW_PROP_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_MBW_PROP_s, MPAMCFG_MBW_PROP_ns, MPAMCFG_MBW_PROP_rt, and MPAMCFG_MBW_PROP_rl must be separate registers:

- The Secure instance (MPAMCFG_MBW_PROP_s) accesses the memory proportional stride bandwidth partitioning used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_MBW_PROP_ns) accesses the memory proportional stride bandwidth partitioning used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_MBW_PROP_rt) accesses the memory proportional stride bandwidth partitioning used for Root PARTIDs.
- The Realm instance (MPAMCFG_MBW_PROP_rl) accesses the memory proportional stride bandwidth partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_MBW_PROP access the memory proportional stride bandwidth partitioning configuration settings for the bandwidth resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_MBW_PROP access the memory proportional stride bandwidth partitioning configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_MBW_PROP access the memory proportional stride bandwidth partitioning configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_MBW_PROP access the memory proportional stride bandwidth partitioning configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_MBW_PROP can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0500	MPAMCFG_MBW_PROP_s

Accesses to this interface are RW.

MPAMCFG_MBW_PROP can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0500	MPAMCFG_MBW_PROP_ns

Accesses to this interface are RW.

MPAMCFG_MBW_PROP can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0500	MPAMCFG_MBW_PROP_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_MBW_PROP can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0500	MPAMCFG_MBW_PROP_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.13 MPAMCFG_MBW_WINWD, MPAM Memory Bandwidth Partitioning Window Width Configuration Register

The MPAMCFG_MBW_WINWD characteristics are:

Purpose

MPAMCFG_MBW_WINWD is a 32-bit register that shows and sets the value of the window width for the PARTID in MPAMCFG_PART_SEL.

MPAMCFG_MBW_WINWD_s reads and controls the bandwidth control window width for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_WINWD_ns reads and controls the bandwidth control window width for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_WINWD_rt reads and controls the bandwidth control window width for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_WINWD_rl reads and controls the bandwidth control window width for the Real PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

MPAMCFG_MBW_WINWD is read-only if MPAMF_MBW_IDR.WINDWR == 0, and the window width is set by the hardware, even if variable.

MPAMCFG_MBW_WINWD is read/write if MPAMF_MBW_IDR.WINDWR == 1, permitting configuration of the window width for each PARTID independently on hardware that supports this functionality.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_MBW_WINWD is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_MBW_PART == 1. Otherwise, direct accesses to MPAMCFG_MBW_WINWD are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_MBW_WINWD is a 32-bit register.

Field descriptions



Bits [31:24]

Reserved, RES0.

US_INT, bits [23:8]

Window width, integer microseconds.

This field reads (and sets) the integer part of the window width in microseconds for the PARTID selected by MPAMCFG_PART_SEL.

US_FRAC, bits [7:0]

Window width, fractional microseconds.

This field reads (and sets) the fractional part of the window width in microseconds for the PARTID selected by MPAMCFG_PART_SEL.

Accessing the MPAMCFG_MBW_WINWD:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_MBW_WINWD_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_MBW_WINWD_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_MBW_WINWD_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_MBW_WINWD_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_MBW_WINWD_s, MPAMCFG_MBW_WINWD_ns, MPAMCFG_MBW_WINWD_rt, and MPAMCFG_MBW_WINWD_rl must be separate registers:

- The Secure instance (MPAMCFG_MBW_WINWD_s) accesses the window width used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_MBW_WINWD_ns) accesses the window width used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_MBW_WINWD_rt) accesses the window width used for Root PARTIDs.
- The Realm instance (MPAMCFG_MBW_WINWD_rl) accesses the window width used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_MBW_WINWD access the window width configuration settings for the bandwidth resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_MBW_WINWD access the window width configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_MBW_WINWD access the window width configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_MBW_WINWD access the window width configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_MBW_WINWD can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0220	MPAMCFG_MBW_WINWD_s

This interface is accessible as follows:

- When MPAMF_MBW_IDR.WINDWR == 0 accesses to this register are RO.
- When MPAMF_MBW_IDR.WINDWR == 1 accesses to this register are RW.

MPAMCFG_MBW_WINWD can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0220	MPAMCFG_MBW_WINWD_ns

This interface is accessible as follows:

- When MPAMF_MBW_IDR.WINDWR == 0 accesses to this register are RO.
- When MPAMF_MBW_IDR.WINDWR == 1 accesses to this register are RW.

MPAMCFG_MBW_WINWD can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0220	MPAMCFG_MBW_WINWD_rt

This interface is accessible as follows:

- When FEAT_RME is implemented and MPAMF_MBW_IDR.WINDWR == 0 accesses to this register are RO.
- When FEAT_RME is implemented and MPAMF_MBW_IDR.WINDWR == 1 accesses to this register are RW.

MPAMCFG_MBW_WINWD can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0220	MPAMCFG_MBW_WINWD_rl

This interface is accessible as follows:

- When FEAT_RME is implemented and MPAMF_MBW_IDR.WINDWR == 0 accesses to this register are RO.
- When FEAT_RME is implemented and MPAMF_MBW_IDR.WINDWR == 1 accesses to this register are RW.

9.4.14 MPAMCFG_PART_SEL, MPAM Partition Configuration Selection Register

The MPAMCFG_PART_SEL characteristics are:

Purpose

Selects a partition ID to configure.

MPAMCFG_PART_SEL_s selects a Secure PARTID to configure. MPAMCFG_PART_SEL_ns selects a Non-secure PARTID to configure. MPAMCFG_PART_SEL_rt selects a Root PARTID to configure. MPAMCFG_PART_SEL_rl selects a Realm PARTID to configure.

After setting this register with a PARTID, software (usually a hypervisor) can perform a series of accesses to MPAMCFG registers to configure parameters for MPAM resource controls to use when requests have that PARTID.

Configurations

The power domain of MPAMCFG_PART_SEL is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and (MPAMF_IDR.HAS_CCAP_PART == 1, or MPAMF_IDR.HAS_CPOR_PART == 1, or MPAMF_IDR.HAS_MBW_PART == 1, or MPAMF_IDR.HAS_PRI_PART == 1, or MPAMF_IDR.HAS_PARTID_NRW == 1, or (MPAMF_IDR.EXT == 0 and MPAMF_IDR.HAS_IMPL_IDR == 1) or (MPAMF_IDR.EXT == 1, MPAMF_IDR.HAS_IMPL_IDR == 1 and MPAMF_IDR.NO_IMPL_PART == 0)). Otherwise, direct accesses to MPAMCFG_PART_SEL are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_PART_SEL is a 32-bit register.

Field descriptions



Bits [31:28]

Reserved, RES0.

RIS, bits [27:24]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented), MPAMF_IDR.EXT == 1 and MPAMF_IDR.HAS_RIS == 1:

Resource Instance Selector. RIS selects one resource to configure through MPAMCFG registers and describe with MPAMF ID registers.

Otherwise:

Reserved, RES0.

Bits [23:17]

Reserved, RES0.

INTERNAL, bit [16]

Internal PARTID.

If MPAMF_IDR.HAS_PARTID_NRW = 0, this field is RAZ/WI.

If MPAMF_IDR.HAS_PARTID_NRW = 1:

0b0 PARTID_SEL is interpreted as a request PARTID and ignored except for use with MPAMCFG_INTPARTID register access.

0b1 PARTID_SEL is interpreted as an internal PARTID and used for access to MPAMCFG control settings except for MPAMCFG_INTPARTID.

If PARTID narrowing is implemented as indicated by MPAMF_IDR.HAS_PARTID_NRW = 1, when accessing other MPAMCFG registers the value of the MPAMCFG_PART_SEL.INTERNAL bit is checked for these conditions:

- When the MPAMCFG_INTPARTID register is read or written, if the value of MPAMCFG_PART_SEL.INTERNAL is not 0, an Unexpected_INTERNAL error is set in MPAMF_ESR.
- When an MPAMCFG register other than MPAMCFG_INTPARTID is read or written, if the value of MPAMCFG_PART_SEL.INTERNAL is not 1, MPAMF_ESR is set to indicate an intPARTID_Range error.

In either error case listed here, the value returned by a read operation is UNPREDICTABLE, and the control settings are not affected by a write.

PARTID_SEL, bits [15:0]

Selects the partition ID to configure.

Reads and writes to other MPAMCFG registers are indexed by PARTID_SEL and by the NS bit used to access MPAMCFG_PART_SEL to access the configuration for a single partition.

Accessing the MPAMCFG_PART_SEL:

This register is within the MPAM feature page memory frames. In a system that supports Secure and Non-secure memory maps, there must be both Secure and Non-secure MPAM feature pages.

MPAMCFG_PART_SEL_s must only be accessible from the Secure MPAM feature page.

MPAMCFG_PART_SEL_ns must only be accessible from the Non-secure MPAM feature page.

MPAMCFG_PART_SEL_s and MPAMCFG_PART_SEL_ns must be separate registers. The Secure instance (MPAMCFG_PART_SEL_s) accesses the PARTID selector used for Secure PARTIDs, and the Non-secure instance (MPAMCFG_PART_SEL_ns) accesses the PARTID selector used for Non-secure PARTIDs.

MPAMCFG_PART_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0100	MPAMCFG_PART_SEL_s

Accesses to this interface are RW.

MPAMCFG_PART_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0100	MPAMCFG_PART_SEL_ns

Accesses to this interface are RW.

MPAMCFG_PART_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0100	MPAMCFG_PART_SEL_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_PART_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0100	MPAMCFG_PART_SEL_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.4.15 MPAMCFG_PRI, MPAM Priority Partition Configuration Register

The MPAMCFG_PRI characteristics are:

Purpose

Controls the internal and downstream priority of requests attributed to the PARTID selected by MPAMCFG_PART_SEL.

MPAMCFG_PRI_s controls the priorities for the Secure PARTID selected by the Secure instance of MPAMCFG_PART_SEL. MPAMCFG_PRI_ns controls the priorities for the Non-secure PARTID selected by the Non-secure instance of MPAMCFG_PART_SEL. MPAMCFG_PRI_rt controls the priorities for the Root PARTID selected by the Root instance of MPAMCFG_PART_SEL. MPAMCFG_PRI_rl controls the priorities for the Realm PARTID selected by the Realm instance of MPAMCFG_PART_SEL.

If MPAMF_IDR.HAS_RIS is 1, the control settings accessed are those of the resource instance currently selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

Configurations

The power domain of MPAMCFG_PRI is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented and MPAMF_IDR.HAS_PRI_PART == 1. Otherwise, direct accesses to MPAMCFG_PRI are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMCFG_PRI is a 32-bit register.

Field descriptions



DSPRI, bits [31:16]

Downstream priority.

If MPAMF_PRI_IDR.HAS_DSPRI == 0, bits of this field are RES0 as this field is not used.

If MPAMF_PRI_IDR.HAS_DSPRI == 1, this field is a priority value applied to downstream communications from this MSC for transactions of the partition selected by MPAMCFG_PART_SEL.

The implemented width of this field is MPAMF_PRI_IDR.DSPRI_WD bits. If the implemented width is less than the width of this field, the least significant bits are used.

The encoding of priority is 0-as-lowest or 0-as-highest priority according to the value of MPAMF_PRI_IDR.DSPRI_0_IS_LOW.

INTPRI, bits [15:0]

Internal priority.

If MPAMF_PRI_IDR.HAS_INTPRI == 0, bits of this field are RES0 as this field is not used.

If MPAMF_PRI_IDR.HAS_INTPRI == 1, this field is a priority value applied internally inside this MSC for transactions of the partition selected by MPAMCFG_PART_SEL.

The implemented width of this field is MPAMF_PRI_IDR.INTPRI_WD bits. If the implemented width is less than the width of this field, the least significant bits are used.

The encoding of priority is 0-as-lowest or 0-as-highest priority according to the value of MPAMF_PRI_IDR.INTPRI_0_IS_LOW.

Accessing the MPAMCFG_PRI:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMCFG_PRI_s must only be accessible from the Secure MPAM feature page.
- MPAMCFG_PRI_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMCFG_PRI_rt must only be accessible from the Root MPAM feature page.
- MPAMCFG_PRI_rl must only be accessible from the Realm MPAM feature page.

MPAMCFG_PRI_s, MPAMCFG_PRI_ns, MPAMCFG_PRI_rt, and MPAMCFG_PRI_rl must be separate registers:

- The Secure instance (MPAMCFG_PRI_s) accesses the priority partitioning used for Secure PARTIDs.
- The Non-secure instance (MPAMCFG_PRI_ns) accesses the priority partitioning used for Non-secure PARTIDs.
- The Root instance (MPAMCFG_PRI_rt) accesses the priority partitioning used for Root PARTIDs.
- The Realm instance (MPAMCFG_PRI_rl) accesses the priority partitioning used for Realm PARTIDs.

When RIS is implemented, loads and stores to MPAMCFG_PRI access the priority partitioning configuration settings for the priority resource instance selected by MPAMCFG_PART_SEL.RIS and the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When RIS is not implemented, loads and stores to MPAMCFG_PRI access the priority partitioning configuration settings for the PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL.

When PARTID narrowing is implemented, loads and stores to MPAMCFG_PRI access the priority partitioning configuration settings for the internal PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 1.

When PARTID narrowing is not implemented, loads and stores to MPAMCFG_PRI access the priority partitioning configuration settings for the request PARTID selected by MPAMCFG_PART_SEL.PARTID_SEL, and MPAMCFG_PART_SEL.INTERNAL must be 0.

MPAMCFG_PRI can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0400	MPAMCFG_PRI_s

Accesses to this interface are RW.

MPAMCFG_PRI can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0400	MPAMCFG_PRI_ns

Accesses to this interface are RW.

MPAMCFG_PRI can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0400	MPAMCFG_PRI_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMCFG_PRI can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0400	MPAMCFG_PRI_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5 Memory-mapped monitoring configuration registers

This section lists the external monitoring configuration registers.

9.5.1 MSMON_CAPT_EVNT, MPAM Capture Event Generation Register

The MSMON_CAPT_EVNT characteristics are:

Purpose

Generates a local capture event when written with bit[0] as 1.

MSMON_CAPT_EVNT_s generates local capture events for Secure monitor instances only or for Secure and Non-secure monitor instances. MSMON_CAPT_EVNT_ns generates local capture events for Non-secure monitor instances only. MSMON_CAPT_EVNT_rt generates local capture events for Root monitor instances only or for Root, Secure, Realm, and Non-secure monitor instances. MSMON_CAPT_EVNT_rl generates local capture events for Realm monitor instances or for for Realm monitor instances or Realm and Non-secure monitor instances.

Configurations

The power domain of MSMON_CAPT_EVNT is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MSMON == 1 and MPAMF_MSMON_IDR.HAS_LOCAL_CAPT_EVNT == 1. Otherwise, direct accesses to MSMON_CAPT_EVNT are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CAPT_EVNT is a 32-bit register.

Field descriptions



Bits [31:2]

Reserved, RES0.

ALL, bit [1]

In the Secure instance of this register:

- If ALL is written as 1 and NOW is also written as 1, signal a capture event to Secure and Non-secure monitor instances in this MSC that are configured with CAPT_EVNT = 7.
- If ALL is written as 0 and NOW is written as 1, signal a capture event to Secure monitor instances in this MSC that are configured with CAPT_EVNT = 7.

In the Non-secure instance of this register, this bit is RAZ/WI.

In the Root instance of this register:

- If ALL is written as 1 and NOW is also written as 1, signal a capture event to Root, Realm, Secure, and Non-secure monitor instances in this MSC that are configured with CAPT_EVNT = 7.
- If ALL is written as 0 and NOW is written as 1, signal a capture event to Root monitor instances within this MSC that are configured with CAPT_EVNT = 7.

In the Realm instance of this register:

- If ALL is written as 1 and NOW is also written as 1, signal a capture event to Realm and Non-secure monitor instances in this MSC that are configured with CAPT_EVNT = 7.
- If ALL is written as 0 and NOW is written as 1, signal a capture event to Realm monitor instances within this MSC that are configured with CAPT_EVNT = 7.

This bit always reads as zero.

0b0 Send capture event only to monitor instances in the same MPAM feature page as this register.

0b1 Send capture event to monitor instances in certain MPAM feature pages as described in the ALL field of this register.

NOW, bit [0]

When written as 1, this bit causes an event to those monitor instances described in the ALL field that have CAPT_EVNT set to the value of 7.

When this bit is written as 0, no event is signaled.

This bit always reads as zero.

Accessing the MSMON_CAPT_EVNT:

This register is within the MPAM feature page memory frames. In a system that supports Secure and Non-secure memory maps, there must be both Secure and Non-secure MPAM feature pages.

MSMON_CAPT_EVNT_s must only be accessible from the Secure MPAM feature page.

MSMON_CAPT_EVNT_ns must only be accessible from the Non-secure MPAM feature page.

MSMON_CAPT_EVNT_s and MSMON_CAPT_EVNT_ns must be separate registers. The Secure instance (MSMON_CAPT_EVNT_s) can generate local capture events for Secure monitor instances only or for Secure and Non-secure monitor instances, and the Non-secure instance (MSMON_CAPT_EVNT_ns) can generate local capture events for Non-secure monitor instances only.

MSMON_CAPT_EVNT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0808	MSMON_CAPT_EVNT_s

Accesses to this interface are WO/RAZ.

MSMON_CAPT_EVNT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0808	MSMON_CAPT_EVNT_ns

Accesses to this interface are WO/RAZ.

MSMON_CAPT_EVNT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0808	MSMON_CAPT_EVNT_rt

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

MSMON_CAPT_EVNT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0808	MSMON_CAPT_EVNT_rl

When FEAT_RME is implemented, accesses to this interface are WO/RAZ.

9.5.2 MSMON_CFG_CSU_CTL, MPAM Memory System Monitor Configure Cache Storage Usage Monitor Control Register

The MSMON_CFG_CSU_CTL characteristics are:

Purpose

Controls the CSU monitor selected by `MSMON_CFG_MON_SEL`.

`MSMON_CFG_CSU_CTL_s` controls the Secure cache storage usage monitor instance selected by the Secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_CSU_CTL_ns` controls Non-secure cache storage usage monitor instance selected by the Non-secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_CSU_CTL_rt` controls the monitor configuration for the Root PARTID selected by the Root instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_CSU_CTL_rl` controls the monitor configuration for the Realm PARTID selected by the Realm instance of `MSMON_CFG_MON_SEL`.

If `MPAMF_IDR.HAS_RIS` is 1, the monitor instance configuration accessed is for the resource instance currently selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance of that resource instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

Configurations

The power domain of `MSMON_CFG_CSU_CTL` is IMPLEMENTATION DEFINED.


This register is present only when `FEAT_MPAM` is implemented, `MPAMF_IDR.HAS_MSMON == 1` and `MPAMF_MSMON_IDR.MSMON_CSU == 1`. Otherwise, direct accesses to `MSMON_CFG_CSU_CTL` are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

`MSMON_CFG_CSU_CTL` is a 32-bit register.

Field descriptions



EN, bit [31]

Enabled.

`0b0` The monitor instance is disabled and must not collect any information.

`0b1` The monitor instance is enabled to collect information according to the configuration of the instance.

CAPT_EVNT, bits [30:28]

Capture event selector.

Select the event that triggers capture from the following:

`0b000` No capture event is triggered.

`0b001` External capture event 1 (optional, but recommended)

`0b010` External capture event 2 (optional)

`0b011` External capture event 3 (optional)

`0b100` External capture event 4 (optional)

`0b101` External capture event 5 (optional)

- 0b110 External capture event 6 (optional)
- 0b111 Capture occurs when a `MSMON_CAPT_EVNT` register in this MSC is written and causes a capture event for the Security state of this monitor. (optional)

The values marked as optional indicate capture event sources that can be omitted in an implementation. Those values representing non-implemented event sources must not trigger a capture event.

When `MPAMF_CSUMON_IDR.HAS_CAPTURE == 0`, access to this field is RAZ/WI.

CAPT_RESET, bit [27]

Reset after capture.

Controls whether the value of `MSMON_CSU` is reset to zero immediately after being copied to `MSMON_CSU_CAPTURE`.

- 0b0 Monitor is not reset on capture.
- 0b1 Monitor is reset on capture.

Because the CSU monitor type produces a measurement rather than a count, it might not make sense to ever reset the value after a capture. If there is no reason to ever reset a CSU monitor, this field is RAZ/WI.

When `MPAMF_CSUMON_IDR.HAS_CAPTURE == 0`, access to this field is RAZ/WI.

OFLOW_STATUS, bit [26]

Overflow status.

Indicates whether the value of `MSMON_CSU` has overflowed.

If `MPAMF_CSUMON_IDR.HAS_CEVRT_OFLW` is 1 or `MPAMF_CSUMON_IDR.HAS_OFLOW_LNKG` is 1, then a store to `MSMON_CSU` when this field is 1 resets this field to 0.

- 0b0 No overflow has occurred.
- 0b1 At least one overflow has occurred since this bit was last written to zero.

If overflow is not possible for a CSU monitor in the implementation, this field is RAZ/WI.

OFLOW_INTR, bit [25]

Overflow Interrupt.

Controls whether an overflow interrupt is generated when the value of `MSMON_CSU` has overflowed.

- 0b0 No interrupt is signaled on an overflow of `MSMON_CSU`.
- 0b1 On overflow, an implementation-specific interrupt is signaled.

If `OFLOW_INTR` is not supported by the implementation, this field is RAZ/WI.

OFLOW_FRZ, bit [24]

Freeze Monitor on Overflow.

Controls whether the value of `MSMON_CSU.VALUE` freezes on an overflow.

- 0b0 Monitor count wraps on overflow.
- 0b1 Monitor count freezes on overflow. The frozen value might be 0 or another value if the monitor overflowed with an increment larger than 1.

If overflow is not possible for a CSU monitor in the implementation, this field is RAZ/WI.

When a `MSMON_CSU.VALUE` of a monitor instance is frozen it does not change until `MSMON_CSU` register for that instance has been written.

OFLOW_CAPT, bit [23]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CSUMON_IDR.HAS_OFLOW_CAPT == 1:

Capture Monitor on Overflow.

- 0b0 Monitor is not captured on an overflow or when affected by an overflow linkage event.
- 0b1 Monitor is captured and the `MSMON_CSU`.{NRDY, VALUE} fields are copied to the monitor instance's capture register on an overflow or when affected by an overflow linkage event. The monitor instance treats an overflow of this monitor instance as a private capture event. If `MSMON_CFG_MBWU_CTL.CEVNT_OFLW` is 1, this monitor instance also treats an overflow linkage event as a capture event.
- If the `OFLOW_FRZ` field is 1, the monitor does not continue to count after the overflow or overflow linkage event. If the `CAPT_RESET` field is 1, the monitor instance resets to 0.

Otherwise:

Reserved, RES0.

SUBTYPE, bits [22:20]

Subtype. Type of cache storage usage counted by this monitor.

This field is not currently used for CSU monitors, but reserved for future use.

This field is RAZ/WI.

Bit [19]

Reserved, RES0.

CEVNT_OFLW, bit [18]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CSUMON_IDR.HAS_CEVNT_OFLW == 1:

Capture Event performs overflow behavior.

- 0b0 On a capture event matching the `CAPT_EVNT` field, the capture behaviors are performed.
- The `MSMON_CSU`.{VALUE, NRDY} fields are transferred to the monitor instance's capture register.
- 0b1 On a capture event matching the `CAPT_EVNT` field, the monitor instance treats a capture event as an overflow and the overflow behaviors are performed.
- The behavior is controlled by the `MSMON_CFG_CSU_CTL`.{OFLOW_FRZ, OFLOW_CAPT, CAPT_RESET} fields. The `MSMON_CFG_CSU_CTL.OFLOW_STATUS` field is set for this monitor instance.

Otherwise:

Reserved, RES0.

MATCH_PMG, bit [17]

Match PMG.

Controls whether the monitor measures only storage used with PMG matching `MSMON_CFG_CSU_FLT.PMG`.

- 0b0 The monitor measures storage used with any PMG value.
- 0b1 The monitor only measures storage used with the PMG value matching `MSMON_CFG_CSU_FLT.PMG`.

If `MATCH_PMG` is 1 and `MATCH_PARTID` is 0, it is CONSTRAINED UNPREDICTABLE whether the monitor instance:

- Measures the storage used with matching PMG and with any PARTID.
- Measures no storage usage, that is, `MSMON_CSU.VALUE` is zero.

- Measures the storage used with matching PMG and PARTID, that is, treats MATCH_PARTID as == 1.

MATCH_PARTID, bit [16]

Match PARTID.

Controls whether the monitor measures only storage used with PARTID matching MSMON_CFG_CSU_FLT.PARTID.

- 0b0 The monitor measures storage used with any PARTID value.
- 0b1 The monitor only measures storage used with the PARTID value matching MSMON_CFG_CSU_FLT.PARTID.

Bits [15:11]

Reserved, RES0.

OFLOW_LNKG, bits [10:8]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CSUMON_IDR.HAS_OFLOW_LNKG == 1:

Overflow linkage event.

Controls signaling of a capture event on overflow of this monitor instance.

- 0b000 Overflow of the monitor instance only affects this monitor instance.
- 0b001 Overflow of this monitor instance signals Capture Event 1.
- 0b010 Overflow of this monitor instance signals Capture Event 2.
- 0b011 Overflow of this monitor instance signals Capture Event 3.
- 0b100 Overflow of this monitor instance signals Capture Event 4.
- 0b101 Overflow of this monitor instance signals Capture Event 5.
- 0b110 Overflow of this monitor instance signals Capture Event 6.
- 0b111 Reserved.

Otherwise:

Reserved, RES0.

TYPE, bits [7:0]

Monitor Type Code. The CSU monitor is TYPE = 0x43.

TYPE is a read-only constant indicating the type of the monitor.

Reads as 0x43.

Access to this field is RO.

Accessing the MSMON_CFG_CSU_CTL:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_CFG_CSU_CTL_s must only be accessible from the Secure MPAM feature page.
- MSMON_CFG_CSU_CTL_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_CFG_CSU_CTL_rt must only be accessible from the Root MPAM feature page.
- MSMON_CFG_CSU_CTL_rl must only be accessible from the Realm MPAM feature page.

MSMON_CFG_CSU_CTL_s, MSMON_CFG_CSU_CTL_ns, MSMON_CFG_CSU_CTL_rt, and MSMON_CFG_CSU_CTL_rl must be separate registers:

- The Secure instance (MSMON_CFG_CSU_CTL_s) accesses the cache storage usage monitor controls used for Secure PARTIDs.
- The Non-secure instance (MSMON_CFG_CSU_CTL_ns) accesses the cache storage usage monitor controls used for Non-secure PARTIDs.
- The Root instance (MSMON_CFG_CSU_CTL_rt) accesses the cache storage usage monitor controls used for Root PARTIDs.
- The Realm instance (MSMON_CFG_CSU_CTL_rl) accesses the cache storage usage monitor controls used for Realm PARTIDs.

When RIS is implemented, loads and stores to MSMON_CFG_CSU_CTL access the cache storage usage monitor configuration settings for the cache resource instance selected by [MSMON_CFG_MON_SEL.RIS](#) and the cache storage usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

When RIS is not implemented, loads and stores to MSMON_CFG_CSU_CTL access the cache storage usage monitor configuration settings for the cache storage usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

MSMON_CFG_CSU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0818	MSMON_CFG_CSU_CTL_s

Accesses to this interface are RW.

MSMON_CFG_CSU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0818	MSMON_CFG_CSU_CTL_ns

Accesses to this interface are RW.

MSMON_CFG_CSU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0818	MSMON_CFG_CSU_CTL_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CFG_CSU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0818	MSMON_CFG_CSU_CTL_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.3 MSMON_CFG_CSU_FLT, MPAM Memory System Monitor Configure Cache Storage Usage Monitor Filter Register

The MSMON_CFG_CSU_FLT characteristics are:

Purpose

Configures PARTID and PMG to measure or count in the CSU monitor selected by [MSMON_CFG_MON_SEL](#).

MSMON_CFG_CSU_FLT_s sets filter conditions for the Secure cache storage usage monitor instance selected by the Secure instance of [MSMON_CFG_MON_SEL](#).

MSMON_CFG_CSU_CTL_ns sets filter conditions for the Non-secure cache storage usage monitor instance selected by the Non-secure instance of [MSMON_CFG_MON_SEL](#).

MSMON_CFG_CSU_FLT_rt sets the filter conditions for the Root PARTID selected by the Root instance of [MSMON_CFG_MON_SEL](#). MSMON_CFG_CSU_FLT_rl sets the filter conditions for the Realm PARTID selected by the Realm instance of [MSMON_CFG_MON_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the monitor instance filter configuration accessed is for the resource instance currently selected by [MSMON_CFG_MON_SEL.RIS](#) and the monitor instance of that resource instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

Configurations

The power domain of MSMON_CFG_CSU_FLT is IMPLEMENTATION DEFINED.


This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MSMON](#) == 1 and [MPAMF_MSMON_IDR.MSMON_CSU](#) == 1. Otherwise, direct accesses to MSMON_CFG_CSU_FLT are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CFG_CSU_FLT is a 32-bit register.

Field descriptions



XCL, bit [31]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_CSUMON_IDR.HAS_XCL == 1:

Exclude Clean. The monitor instance does not count cache storage used by lines in an unmodified cache state.

0b0 Monitor instance counts cache storage in modified and unmodified cache lines.

0b1 Monitor instance counts cache storage in modified cache lines only.

Otherwise:

Reserved, RES0.

Bits [30:24]

Reserved, RES0.

PMG, bits [23:16]

Performance monitoring group to filter cache storage usage monitoring.

If [MSMON_CFG_CSU_CTL.MATCH_PMG](#) is 0, this field is not used to match cache storage to a PMG and the contents of this field is ignored.

If `MSMON_CFG_CSU_CTL.MATCH_PMG` is 1 and `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 1, the monitor instance selected by `MSMON_CFG_MON_SEL` measures or counts cache storage labeled with `PMG` equal to this field and `PARTID` equal to the `PARTID` field.

If `MSMON_CFG_CSU_CTL.MATCH_PMG` is 1 and `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 0, the behavior of the monitor instance selected by `MSMON_CFG_MON_SEL` is `CONSTRAINED UNPREDICTABLE`. See `MSMON_CFG_CSU_CTL.MATCH_PMG` for more information.

PARTID, bits [15:0]

Partition ID to filter cache storage usage monitoring.

If `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 0 and `MSMON_CFG_CSU_CTL.MATCH_PMG` is 0, the monitor measures all allocated cache storage.

If `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 0 and `MSMON_CFG_CSU_CTL.MATCH_PMG` is 1, the behavior of the monitor is `CONSTRAINED UNPREDICTABLE`. See the description of `MSMON_CFG_CSU_CTL.MATCH_PMG`.

If `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 1 and `MSMON_CFG_CSU_CTL.MATCH_PMG` is 0, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts cache storage labeled with `PARTID` equal to this field.

If `MSMON_CFG_CSU_CTL.MATCH_PARTID` is 1 and `MSMON_CFG_CSU_CTL.MATCH_PMG` is 1, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts cache storage labeled with `PARTID` equal to this field and `PMG` equal to the `PMG` field.

Accessing the MSMON_CFG_CSU_FLT:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MSMON_CFG_CSU_FLT_s` must only be accessible from the Secure MPAM feature page.
- `MSMON_CFG_CSU_FLT_ns` must only be accessible from the Non-secure MPAM feature page.
- `MSMON_CFG_CSU_FLT_rt` must only be accessible from the Root MPAM feature page.
- `MSMON_CFG_CSU_FLT_rl` must only be accessible from the Realm MPAM feature page.

`MSMON_CFG_CSU_FLT_s`, `MSMON_CFG_CSU_FLT_ns`, `MSMON_CFG_CSU_FLT_rt`, and `MSMON_CFG_CSU_FLT_rl` must be separate registers:

- The Secure instance (`MSMON_CFG_CSU_FLT_s`) accesses the `PARTID` and `PMG` matching for a cache storage usage monitor used for Secure `PARTIDs`.
- The Non-secure instance (`MSMON_CFG_CSU_FLT_ns`) accesses the `PARTID` and `PMG` matching for a cache storage usage monitor used for Non-secure `PARTIDs`.
- The Root instance (`MSMON_CFG_CSU_FLT_rt`) accesses the `PARTID` and `PMG` matching for a cache storage usage monitor used for Root `PARTIDs`.
- The Realm instance (`MSMON_CFG_CSU_FLT_rl`) accesses the `PARTID` and `PMG` matching for a cache storage usage monitor used for Realm `PARTIDs`.

When `RIS` is implemented, loads and stores to `MSMON_CFG_CSU_FLT` access the monitor configuration settings for the resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the cache storage usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When `RIS` is not implemented, loads and stores to `MSMON_CFG_CSU_FLT` access the monitor configuration settings for the cache storage usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

MSMON_CFG_CSU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0810	MSMON_CFG_CSU_FLT_s

Accesses to this interface are RW.

MSMON_CFG_CSU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0810	MSMON_CFG_CSU_FLT_ns

Accesses to this interface are RW.

MSMON_CFG_CSU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0810	MSMON_CFG_CSU_FLT_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CFG_CSU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0810	MSMON_CFG_CSU_FLT_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.4 MSMON_CFG_MBWU_CTL, MPAM Memory System Monitor Configure Memory Bandwidth Usage Monitor Control Register

The MSMON_CFG_MBWU_CTL characteristics are:

Purpose

Controls the MBWU monitor selected by `MSMON_CFG_MON_SEL`.

`MSMON_CFG_MBWU_CTL_s` controls the Secure memory bandwidth usage monitor instance selected by the Secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_MBWU_CTL_ns` controls Non-secure memory bandwidth usage monitor instance selected by the Non-secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_MBWU_CTL_rt` controls the monitor configuration for the Root PARTID selected by the Root instance of `MSMON_CFG_MON_SEL`. `MSMON_CFG_MBWU_CTL_rl` controls the monitor configuration for the Realm PARTID selected by the Realm instance of `MSMON_CFG_MON_SEL`.

If `MPAMF_IDR.HAS_RIS` is 1, the monitor instance configuration accessed is for the resource instance currently selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance of that resource instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

Configurations

The power domain of `MSMON_CFG_MBWU_CTL` is IMPLEMENTATION DEFINED.


This register is present only when `FEAT_MPAM` is implemented, `MPAMF_IDR.HAS_MSMON == 1` and `MPAMF_MSMON_IDR.MSMON_MBWU == 1`. Otherwise, direct accesses to `MSMON_CFG_MBWU_CTL` are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

`MSMON_CFG_MBWU_CTL` is a 32-bit register.

Field descriptions



EN, bit [31]

Enabled.

`0b0` The monitor instance is disabled and must not collect any information.

`0b1` The monitor instance is enabled to collect information according to the configuration of the instance.

CAPT_EVNT, bits [30:28]

Capture event selector.

When the selected capture event occurs, `MSMON_MBWU` of the monitor instance is copied to `MSMON_MBWU_CAPTURE` of the same instance. If the long counter is also implemented, `MSMON_MBWU_L` is also copied to `MSMON_MBWU_L_CAPTURE`.

Select the event that triggers capture from the following:

`0b000` No capture event is triggered.

`0b001` External capture event 1 (optional, but recommended)

`0b010` External capture event 2 (optional)

0b011	External capture event 3 (optional)
0b100	External capture event 4 (optional)
0b101	External capture event 5 (optional)
0b110	External capture event 6 (optional)
0b111	Capture occurs when a <code>MSMON_CAPT_EVNT</code> register in this MSC is written and causes a capture event for the Security state of this monitor. (optional)

The values marked as optional indicate capture event sources that can be omitted in an implementation. Those values representing non-implemented event sources must not trigger a capture event.

If capture is not implemented for the MBWU monitor type as indicated by `MPAMF_MBWUMON_IDR.HAS_CAPTURE = 0`, this field is RAZ/WI.

CAPT_RESET, bit [27]

Reset `MSMON_MBWU.VALUE` after capture.

Controls whether the `VALUE` field of the monitor instance is reset to zero immediately after being copied to the corresponding capture register.

0b0	<code>MSMON_MBWU.VALUE</code> field of the monitor instance is not reset on capture.
0b1	<code>MSMON_MBWU.VALUE</code> field of the monitor instance is reset on capture.

If capture is not implemented for the MBWU monitor type as indicated by `MPAMF_MBWUMON_IDR.HAS_CAPTURE = 0`, this field is RAZ/WI.

This control bit affects both `MSMON_MBWU` and `MSMON_MBWU_L` in implementations that include `MSMON_MBWU_L`.

OFLOW_STATUS, bit [26]

Overflow status.

Indicates whether the value of `MSMON_MBWU` has overflowed.

0b0	<code>MSMON_MBWU.VALUE</code> has not overflowed.
0b1	<code>MSMON_MBWU.VALUE</code> has overflowed at least once since this bit was last written to zero.

If overflow is not possible for an MBWU monitor in the MSC implementation, this field is RAZ/WI.

Overflow status for `MSMON_MBWU_L.VALUE` is reported in `MSMON_CFG_MBWU_CTL.OFLOW_STATUS_L`.

If `MPAMF_MBWUMON_IDR.HAS_CEVTN_OFLW` is 1 or `MPAMF_MBWUMON_IDR.HAS_OFLOW_LNKG` is 1, then a store to `MSMON_MBWU` when this field is 1 resets this field to 0.

OFLOW_INTR, bit [25]

Enable interrupt on overflow of `MSMON_MBWU.VALUE`.

0b0	No interrupt is signaled on an overflow of <code>MSMON_MBWU.VALUE</code> .
0b1	An implementation-specific interrupt is signaled on an overflow of <code>MSMON_MBWU.VALUE</code> .

If overflow is not possible for an MBWU monitor in the MSC implementation, this field is RAZ/WI.

If overflow interrupt is not supported by the MSC implementation, this field is RAZ/WI.

Interrupt enable for overflow of `MSMON_MBWU_L.VALUE` is controlled by `MSMON_CFG_MBWU_CTL.OFLOW_INTR_L`.

OFLOW_FRZ, bit [24]

Freeze monitor instance on overflow.

Controls whether `MSMON_MBWU.VALUE` field of the monitor instance freezes on an overflow.

0b0	<code>MSMON_MBWU.VALUE</code> field of the monitor instance wraps on overflow.
-----	--

0b1 [MSMON_MBWU.VALUE](#) field of the monitor instance freezes on overflow. If the increment that caused the overflow was 1, the frozen value is the post-increment value of 0. If the increment that caused the overflow was larger than 1, the frozen value of the monitor might be 0 or a larger value less than the final increment.

If overflow is not possible for the instance of the MBWU monitor in the implementation, this field is RAZ/WI.

When a [MSMON_MBWU.VALUE](#) of a monitor instance is frozen it does not change until [MSMON_CSU](#) register for that instance has been written. If the monitor implements both [MSMON_MBWU](#) and [MSMON_MBWU_L](#) registers, both are frozen. A write to a frozen register unfreezes the count for just that register.

This control bit affects both [MSMON_MBWU](#) and [MSMON_MBWU_L](#) in implementations that include [MSMON_MBWU_L](#).

OFLOW_CAPT, bit [23]

When ([FEAT_MPAMv0p1](#) is implemented or [FEAT_MPAMv1p1](#) is implemented) and [MPAMF_MBWUMON_IDR.HAS_OFLOW_CAPT](#) == 1:

Capture Monitor on Overflow.

0b0 Monitor register [MSMON_MBWU](#) is not captured on an overflow or when affected by an overflow linkage event.

0b1 Monitor register [MSMON_MBWU](#) is captured and the [MSMON_MBWU](#).{NRDY, VALUE} fields are copied to the monitor instance's [MSMON_MBWU_CAPTURE](#) register on an overflow or when affected by an overflow linkage event. The monitor instance treats an overflow of this monitor instance as a private capture event. If [MSMON_CFG_MBWU_CTL.CEVNT_OFLW](#) is 1, this monitor instance also treats an overflow linkage event as a capture event.

If [OFLOW_FRZ](#) is 1, the monitor does not continue to count after the overflow or overflow linkage event. If [CAPT_RESET](#) is 1, the monitor instance resets to 0.

This bit does not control whether [MSMON_MBWU_L](#) is captured on an overflow or overflow linkage event. See [MSMON_CFG_MBWU_CTL.OFLOW_CAPT_L](#).

Otherwise:

Reserved, RES0.

SUBTYPE, bits [22:20]

Subtype. Type of bandwidth counted by this monitor.

This field is not currently used for MBWU monitors, but reserved for future use.

This field is RAZ/WI.

SCLEN, bit [19]

[MSMON_MBWU.VALUE](#) Scaling Enable.

Enables scaling of [MSMON_MBWU.VALUE](#) by [MPAMF_MBWUMON_IDR.SCALE](#).

0b0 [MSMON_MBWU.VALUE](#) has bytes counted by the monitor instance.

0b1 [MSMON_MBWU.VALUE](#) has bytes counted by the monitor instance, shifted right by [MPAMF_MBWUMON_IDR.SCALE](#).

CEVNT_OFLW, bit [18]

When ([FEAT_MPAMv0p1](#) is implemented or [FEAT_MPAMv1p1](#) is implemented) and [MPAMF_MBWUMON_IDR.HAS_CEVNT_OFLW](#) == 1:

Capture Event performs overflow behavior.

0b0 On a capture event matching the [CAPT_EVNT](#) field the capture behaviors are performed.

The NRDY and VALUE fields are transferred to the monitor instance's capture register.

0b1 On a capture event matching the [CAPT_EVNT](#) field the monitor instance treats a capture event as an overflow and the overflow behaviors are performed.

The behavior is controlled by the `MSMON_CFG_MBWU_CTL`.`{OFLOW_FRZ, OFLOW_CAPT, OFLOW_CAPT_L, CAPT_RESET}` fields. The `MSMON_CFG_MBWU_CTL`.`{OFLOW_STATUS, OFLOW_STATUS_L}` fields are set for this monitor instance.

Otherwise:

Reserved, RES0.

MATCH_PMG, bit [17]

Match PMG.

Controls whether the monitor instance only counts data transferred with PMG matching `MSMON_CFG_MBWU_FLT.PMG`.

- 0b0 The monitor instance counts data transferred with any PMG value.
- 0b1 The monitor instance only counts data transferred with the PMG value matching `MSMON_CFG_MBWU_FLT.PMG`.

MATCH_PARTID, bit [16]

Match PARTID.

Controls whether the monitor instance counts only data transferred with PARTID matching `MSMON_CFG_MBWU_FLT.PARTID`.

- 0b0 The monitor instance counts data transferred with any PARTID value.
- 0b1 The monitor instance only counts data transferred with the PARTID value matching `MSMON_CFG_MBWU_FLT.PARTID`.

OFLOW_STATUS_L, bit [15]

When `FEAT_MPAMv0p1` is implemented or `FEAT_MPAMv1p1` is implemented:

Overflow Status of `MSMON_MBWU_L.VALUE` of the monitor instance.

Indicates whether `MSMON_MBWU_L.VALUE` has overflowed.

- 0b0 `MSMON_MBWU_L.VALUE` has not overflowed.
- 0b1 `MSMON_MBWU_L.VALUE` has overflowed at least once since this bit was last written to zero.

If `MPAMF_MBWUMON_IDR.HAS_LONG == 0`, this bit is RES0.

Overflow status of `MSMON_MBWU.VALUE` is reported in `MSMON_CFG_MBWU_CTL.OFLOW_STATUS`.

If `MPAMF_MBWUMON_IDR.HAS_CEVRT_OFLW` is 1 or `MPAMF_MBWUMON_IDR.HAS_OFLOW_LNKG` is 1, then a store to `MSMON_MBWU_L` when this field is 1 resets this field to 0.

Otherwise:

Reserved, RES0.

OFLOW_INTR_L, bit [14]

When (`FEAT_MPAMv0p1` is implemented or `FEAT_MPAMv1p1` is implemented) and `MPAMF_MBWUMON_IDR.HAS_LONG == 1`:

Overflow Interrupt for `MSMON_MBWU_L`.

Controls whether an MPAM overflow interrupt is generated when `MSMON_MBWU_L.VALUE` overflows.

- 0b0 No interrupt is signaled on an overflow of `MSMON_MBWU_L.VALUE`.
- 0b1 An implementation-specific interrupt is signaled on overflow of `MSMON_MBWU_L.VALUE`.

If overflow is not possible for an MBWU monitor in the MSC implementation, this field is RAZ/WI.

If the overflow interrupt is not supported by the MSC implementation, this field is RAZ/WI.

Otherwise:

Reserved, RES0.

OFLOW_CAPT_L, bit [13]

**When (*FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented*),
MPAMF_MBWUMON_IDR.HAS_LONG == 1 and
*MPAMF_MBWUMON_IDR.HAS_OFLOW_CAPT == 1:***

Capture Long Monitor on Overflow.

Controls whether `MSMON_MBWU_L` is copied to `MSMON_MBWU_L_CAPTURE` on an overflow or an overflow linkage event.

0b0 Monitor register `MSMON_MBWU_L` is not captured on an overflow or when affected by an overflow linkage event.

0b1 Monitor register `MSMON_MBWU_L` is captured on an overflow or when affected by an overflow linkage event. If `OFLOW_FRZ` is 1, the monitor does not continue to count after the overflow or overflow linkage event. If `CAPT_RESET` is 1, the monitor instance resets to 0.

If this bit is 1, this monitor instance treats an overflow of this monitor instance as a private capture event.

If this bit is 1, this monitor instance also treats overflow linkage events for which it qualifies as a private capture event.

Otherwise:

Reserved, RES0.

Bits [12:11]

Reserved, RES0.

OFLOW_LNKG, bits [10:8]

**When (*FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented*) and
*MPAMF_MBWUMON_IDR.HAS_OFLOW_LNKG == 1:***

Overflow linkage event.

Controls signaling of a capture event on overflow of this monitor instance.

0b000 Overflow of the monitor instance only affects this monitor instance.

0b001 Overflow of this monitor instance signals Capture Event 1.

0b010 Overflow of this monitor instance signals Capture Event 2.

0b011 Overflow of this monitor instance signals Capture Event 3.

0b100 Overflow of this monitor instance signals Capture Event 4.

0b101 Overflow of this monitor instance signals Capture Event 5.

0b110 Overflow of this monitor instance signals Capture Event 6.

0b111 Reserved.

Otherwise:

Reserved, RES0.

TYPE, bits [7:0]

Monitor Type Code. The MBWU monitor is `TYPE = 0x42`.

`TYPE` is a read-only constant indicating the type of the monitor.

Reads as `0x42`.

Access to this field is RO.

Accessing the MSMON_CFG_MBWU_CTL:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_CFG_MBWU_CTL_s must only be accessible from the Secure MPAM feature page.
- MSMON_CFG_MBWU_CTL_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_CFG_MBWU_CTL_rt must only be accessible from the Root MPAM feature page.
- MSMON_CFG_MBWU_CTL_rl must only be accessible from the Realm MPAM feature page.

MSMON_CFG_MBWU_CTL_s, MSMON_CFG_MBWU_CTL_ns, MSMON_CFG_MBWU_CTL_rt, and MSMON_CFG_MBWU_CTL_rl must be separate registers:

- The Secure instance (MSMON_CFG_MBWU_CTL_s) accesses the memory bandwidth usage monitor controls used for Secure PARTIDs.
- The Non-secure instance (MSMON_CFG_MBWU_CTL_ns) accesses the memory bandwidth usage monitor controls used for Non-secure PARTIDs.
- The Root instance (MSMON_CFG_MBWU_CTL_rt) accesses the memory bandwidth usage monitor controls used for Root PARTIDs.
- The Realm instance (MSMON_CFG_MBWU_CTL_rl) accesses the memory bandwidth usage monitor controls used for Realm PARTIDs.

When RIS is implemented, loads and stores to MSMON_CFG_MBWU_CTL access the monitor configuration settings for the bandwidth resource instance selected by [MSMON_CFG_MON_SEL.RIS](#) and the memory bandwidth usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

When RIS is not implemented, loads and stores to MSMON_CFG_MBWU_CTL access the monitor configuration settings for the memory bandwidth usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

MSMON_CFG_MBWU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0828	MSMON_CFG_MBWU_CTL_s

Accesses to this interface are RW.

MSMON_CFG_MBWU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0828	MSMON_CFG_MBWU_CTL_ns

Accesses to this interface are RW.

MSMON_CFG_MBWU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0828	MSMON_CFG_MBWU_CTL_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CFG_MBWU_CTL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0828	MSMON_CFG_MBWU_CTL_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.5 MSMON_CFG_MBWU_FLT, MPAM Memory System Monitor Configure Memory Bandwidth Usage Monitor Filter Register

The MSMON_CFG_MBWU_FLT characteristics are:

Purpose

Controls PARTID and PMG to measure or count in the MBWU monitor selected by [MSMON_CFG_MON_SEL](#).

MSMON_CFG_MBWU_FLT_s sets filter conditions for the Secure memory bandwidth usage monitor instance selected by the Secure instance of [MSMON_CFG_MON_SEL](#).

MSMON_CFG_MBWU_CTL_ns sets filter conditions for the Non-secure memory bandwidth usage monitor instance selected by the Non-secure instance of [MSMON_CFG_MON_SEL](#).

MSMON_CFG_CSU_FLT_rt sets the filter conditions for the Root PARTID selected by the Root instance of [MSMON_CFG_MON_SEL](#). MSMON_CFG_CSU_FLT_rl sets the filter conditions for the Realm PARTID selected by the Realm instance of [MSMON_CFG_MON_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the monitor instance filter configuration accessed is for the resource instance currently selected by [MSMON_CFG_MON_SEL.RIS](#) and the monitor instance of that resource instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

Configurations

The power domain of MSMON_CFG_MBWU_FLT is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MSMON](#) == 1 and [MPAMF_MSMON_IDR.MSMON_MBWU](#) == 1. Otherwise, direct accesses to MSMON_CFG_MBWU_FLT are RES0.


The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CFG_MBWU_FLT is a 32-bit register.

Field descriptions

When FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented:



RW filtering.

RWBW, bits [31:30]

When [MPAMF_MBWUMON_IDR.HAS_RWBW](#) == 1:

Read/write bandwidth filter. Configures the selected monitor instance to count all bandwidth, only read bandwidth or only write bandwidth.

- 0b00 Monitor instance counts read bandwidth and write bandwidth.
- 0b01 Monitor instance counts write bandwidth only.
- 0b10 Monitor instance counts read bandwidth only.
- 0b11 Reserved.

Otherwise:

Reserved, RES0.

Bits [29:24]

Reserved, RES0.

PMG, bits [23:16]

Performance monitoring group to filter memory bandwidth usage monitoring.

If `MSMON_CFG_MBWU_CTL.MATCH_PMG == 0`, this field is not used to match memory bandwidth to a PMG and the contents of this field is ignored.

If `MSMON_CFG_MBWU_CTL.MATCH_PMG == 1`, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts memory bandwidth labeled with PMG equal to this field.


PARTID, bits [15:0]

Partition ID to filter memory bandwidth usage monitoring.

If `MSMON_CFG_MBWU_CTL.MATCH_PARTID == 0`, this field is not used to match memory bandwidth to a PARTID and the contents of this field is ignored.

If `MSMON_CFG_MBWU_CTL.MATCH_PARTID == 1`, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts memory bandwidth labeled with PARTID equal to this field.

Otherwise:



Bits [31:24]

Reserved, RES0.

PMG, bits [23:16]

Performance monitoring group to filter memory bandwidth usage monitoring.

If `MSMON_CFG_MBWU_CTL.MATCH_PMG == 0`, this field is not used to match memory bandwidth to a PMG and the contents of this field is ignored.

If `MSMON_CFG_MBWU_CTL.MATCH_PMG == 1`, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts memory bandwidth labeled with PMG equal to this field.

PARTID, bits [15:0]

Partition ID to filter memory bandwidth usage monitoring.

If `MSMON_CFG_MBWU_CTL.MATCH_PARTID == 0`, this field is not used to match memory bandwidth to a PARTID and the contents of this field is ignored.

If `MSMON_CFG_MBWU_CTL.MATCH_PARTID == 1`, the monitor selected by `MSMON_CFG_MON_SEL` measures or counts memory bandwidth labeled with PARTID equal to this field.

Accessing the `MSMON_CFG_MBWU_FLT`:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MSMON_CFG_MBWU_FLT_s` must only be accessible from the Secure MPAM feature page.
- `MSMON_CFG_MBWU_FLT_ns` must only be accessible from the Non-secure MPAM feature page.
- `MSMON_CFG_MBWU_FLT_rt` must only be accessible from the Root MPAM feature page.
- `MSMON_CFG_MBWU_FLT_rl` must only be accessible from the Realm MPAM feature page.

MSMON_CFG_MBWU_FLT_s, MSMON_CFG_MBWU_FLT_ns, MSMON_CFG_MBWU_FLT_rt, and MSMON_CFG_MBWU_FLT_rl must be separate registers:

- The Secure instance (MSMON_CFG_MBWU_FLT_s) accesses the PARTID and PMG matching for a memory bandwidth usage monitor used for Secure PARTIDs.
- The Non-secure instance (MSMON_CFG_MBWU_FLT_ns) accesses the PARTID and PMG matching for a memory bandwidth usage monitor used for Non-secure PARTIDs.
- The Root instance (MSMON_CFG_MBWU_FLT_rt) accesses the PARTID and PMG matching for a memory bandwidth usage monitor used for Root PARTIDs.
- The Realm instance (MSMON_CFG_MBWU_FLT_rl) accesses the PARTID and PMG matching for a memory bandwidth usage monitor used for Realm PARTIDs.

When RIS is implemented, loads and stores to MSMON_CFG_MBWU_FLT access the monitor configuration settings for the bandwidth resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When RIS is not implemented, loads and stores to MSMON_CFG_MBWU_FLT access the monitor configuration settings for the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

MSMON_CFG_MBWU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0820	MSMON_CFG_MBWU_FLT_s

Accesses to this interface are RW.

MSMON_CFG_MBWU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0820	MSMON_CFG_MBWU_FLT_ns

Accesses to this interface are RW.

MSMON_CFG_MBWU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0820	MSMON_CFG_MBWU_FLT_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CFG_MBWU_FLT can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0820	MSMON_CFG_MBWU_FLT_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.6 MSMON_CFG_MON_SEL, MPAM Monitor Instance Selection Register

The MSMON_CFG_MON_SEL characteristics are:

Purpose

Selects a monitor instance to access through the MSMON configuration and counter registers.

MSMON_CFG_MON_SEL_s selects a Secure monitor instance to access via the Secure MPAM feature page. MSMON_CFG_MON_SEL_{ns} selects a Non-secure monitor instance to access via the Non-secure MPAM feature page. MSMON_CFG_MON_SEL_{rt} selects a Root monitor instance to access via the Root MPAM feature page. MSMON_CFG_MON_SEL_{rl} selects a Realm monitor instance to access via the Realm MPAM feature page.

Note

Different performance monitoring features within an MSC could have different numbers of monitor instances. See the NUM_MON field in the corresponding ID register. This means that a monitor out-of-bounds error might be signaled when an MSMON_CFG register is accessed because the value in MSMON_CFG_MON_SEL.MON_SEL is too large for the particular monitoring feature.

To configure a monitor, set MON_SEL in this register to the index of the monitor instance to configure, then write to the MSMON_CFG_x register to set the configuration of the monitor. At a later time, read the monitor register (for example, MSMON_MBWU) to get the value of the monitor.

Configurations

The power domain of MSMON_CFG_MON_SEL is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented and (MPAMF_IDR.HAS_MSMON == 1, or (MPAMF_IDR.HAS_IMPL_IDR == 1 and MPAMF_IDR.EXT == 0) or (MPAMF_IDR.HAS_IMPL_IDR == 1, MPAMF_IDR.EXT == 1 and MPAMF_IDR.NO_IMPL_MSMON == 0)). Otherwise, direct accesses to MSMON_CFG_MON_SEL are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CFG_MON_SEL is a 32-bit register.

Field descriptions



Bits [31:28]

Reserved, RES0.

RIS, bits [27:24]

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented), MPAMF_IDR.EXT == 1 and MPAMF_IDR.HAS_RIS == 1:

Resource Instance Selector. RIS selects one resource to configure through MSMON_CFG registers.

Otherwise:

Reserved, RES0.

Bits [23:16]

Reserved, RES0.

MON_SEL, bits [15:0]

Selects the monitor instance to configure or read.

Reads and writes to other MSMON registers are indexed by MON_SEL and by the NS bit used to access MSMON_CFG_MON_SEL to access the configuration for a single monitor.

Accessing the MSMON_CFG_MON_SEL:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_CFG_MON_SEL_s must only be accessible from the Secure MPAM feature page.
- MSMON_CFG_MON_SEL_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_CFG_MON_SEL_rt must only be accessible from the Root MPAM feature page.
- MSMON_CFG_MON_SEL_rl must only be accessible from the Realm MPAM feature page.

MSMON_CFG_MON_SEL_s, MSMON_CFG_MON_SEL_ns, MSMON_CFG_MON_SEL_rt, and MSMON_CFG_MON_SEL_rl must be separate registers:

- The Secure instance (MSMON_CFG_MON_SEL_s) accesses the monitor instance selector used for Secure PARTIDs.
- The Non-secure instance (MSMON_CFG_MON_SEL_ns) accesses the monitor instance selector used for Non-secure PARTIDs.
- The Root instance (MSMON_CFG_MON_SEL_rt) accesses the monitor instance selector used for Root PARTIDs.
- The Realm instance (MSMON_CFG_MON_SEL_rl) accesses the monitor instance selector used for Realm PARTIDs.

MSMON_CFG_MON_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0800	MSMON_CFG_MON_SEL_s

Accesses to this interface are RW.

MSMON_CFG_MON_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0800	MSMON_CFG_MON_SEL_ns

Accesses to this interface are RW.

MSMON_CFG_MON_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0800	MSMON_CFG_MON_SEL_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CFG_MON_SEL can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0800	MSMON_CFG_MON_SEL_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.7 MSMON_CSU, MPAM Cache Storage Usage Monitor Register

The MSMON_CSU characteristics are:

Purpose

Accesses the CSU monitor instance selected by `MSMON_CFG_MON_SEL`.

`MSMON_CSU_s` is a Secure cache storage usage monitor instance selected by the Secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CSU_ns` is a Non-secure cache storage usage monitor instance selected by the Non-secure instance of `MSMON_CFG_MON_SEL`. `MSMON_CSU_rt` is a Root cache storage usage monitor instance selected by the Root instance of `MSMON_CFG_MON_SEL`. `MSMON_CSU_rl` is a Realm cache storage usage monitor instance selected by the Realm instance of `MSMON_CFG_MON_SEL`.

If `MPAMF_IDR.HAS_RIS` is 1, the monitor instance accessed is for the resource instance currently selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance of that resource instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

Configurations

The power domain of `MSMON_CSU` is IMPLEMENTATION DEFINED.

This register is present only when `FEAT_MPAM` is implemented, `MPAMF_IDR.HAS_MSMON == 1` and `MPAMF_MSMON_IDR.MSMON_CSU == 1`. Otherwise, direct accesses to `MSMON_CSU` are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

`MSMON_CSU` is a 32-bit register.

Field descriptions



NRDY, bit [31]

Not Ready. Indicates whether the monitor instance has possibly inaccurate data.

0b0 The monitor instance is ready and the `MSMON_CSU.VALUE` field is accurate.

0b1 The monitor instance is not ready and the contents of the `MSMON_CSU.VALUE` field might be inaccurate or otherwise not represent the actual cache storage usage.

VALUE, bits [30:0]

Cache storage usage measurement value if `MSMON_CSU.NRDY` is 0. Invalid if `MSMON_CSU.NRDY` is 1.

`VALUE` is the cache storage usage measured in bytes meeting the criteria set in `MSMON_CFG_CSU_FLT` and `MSMON_CFG_CSU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

Accessing the MSMON_CSU:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MSMON_CSU_s` must only be accessible from the Secure MPAM feature page.
- `MSMON_CSU_ns` must only be accessible from the Non-secure MPAM feature page.

- MSMON_CSU_rt must only be accessible from the Root MPAM feature page.
- MSMON_CSU_rl must only be accessible from the Realm MPAM feature page.

MSMON_CSU_s, MSMON_CSU_ns, MSMON_CSU_rt, and MSMON_CSU_rl must be separate registers:

- The Secure instance (MSMON_CSU_s) accesses the cache storage usage monitor used for Secure PARTIDs.
- The Non-secure instance (MSMON_CSU_ns) accesses the cache storage usage monitor used for Non-secure PARTIDs.
- The Root instance (MSMON_CSU_rt) accesses the cache storage usage monitor used for Root PARTIDs.
- The Realm instance (MSMON_CSU_rl) accesses the cache storage usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to MSMON_CSU access the cache storage usage monitor instance for the cache resource instance selected by [MSMON_CFG_MON_SEL.RIS](#) and the cache storage usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

When RIS is not implemented, reads and writes to MSMON_CSU access the cache storage usage monitor instance for the cache storage usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

MSMON_CSU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0840	MSMON_CSU_s

This interface is accessible as follows:

- When MPAMF_CSUMON_IDR.CSU_RO == 0 accesses to this register are RW.
- When MPAMF_CSUMON_IDR.CSU_RO == 1 accesses to this register are RO.

MSMON_CSU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0840	MSMON_CSU_ns

This interface is accessible as follows:

- When MPAMF_CSUMON_IDR.CSU_RO == 0 accesses to this register are RW.
- When MPAMF_CSUMON_IDR.CSU_RO == 1 accesses to this register are RO.

MSMON_CSU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0840	MSMON_CSU_rt

This interface is accessible as follows:

- When FEAT_RME is implemented and MPAMF_CSUMON_IDR.CSU_RO == 0 accesses to this register are RW.
- When FEAT_RME is implemented and MPAMF_CSUMON_IDR.CSU_RO == 1 accesses to this register are RO.

MSMON_CSU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0840	MSMON_CSU_r1

This interface is accessible as follows:

- When FEAT_RME is implemented and MPAMF_CSUMON_IDR.CSU_RO == 0 accesses to this register are RW.
- When FEAT_RME is implemented and MPAMF_CSUMON_IDR.CSU_RO == 1 accesses to this register are RO.

9.5.8 MSMON_CSU_CAPTURE, MPAM Cache Storage Usage Monitor Capture Register

The MSMON_CSU_CAPTURE characteristics are:

Purpose

MSMON_CSU_CAPTURE is a 32-bit read/write register that accesses the captured [MSMON_CSU](#) monitor instance selected by [MSMON_CFG_MON_SEL](#).

MSMON_CSU_CAPTURE_s is the Secure cache storage usage monitor capture instance selected by the Secure instance of [MSMON_CFG_MON_SEL](#). MSMON_CSU_CAPTURE_ns is the Non-secure cache storage usage monitor capture instance selected by the Non-secure instance of [MSMON_CFG_MON_SEL](#). MSMON_CSU_CAPTURE_rt is a Root cache storage usage monitor capture instance selected by the Root instance of [MSMON_CFG_MON_SEL](#).

MSMON_CSU_CAPTURE_rl is a Realm cache storage usage monitor capture instance selected by the Realm instance of [MSMON_CFG_MON_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the monitor instance capture register accessed is for the resource instance currently selected by [MSMON_CFG_MON_SEL.RIS](#) and the monitor instance of that resource instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

Configurations

The power domain of MSMON_CSU_CAPTURE is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MSMON](#) == 1, [MPAMF_MSMON_IDR.MSMON_CSU](#) == 1 and [MPAMF_CSUMON_IDR.HAS_CAPTURE](#) == 1. Otherwise, direct accesses to MSMON_CSU_CAPTURE are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CSU_CAPTURE is a 32-bit register.

Field descriptions



NRDY, bit [31]

Not Ready. Indicates whether the captured monitor value has possibly inaccurate data.

- 0b0 The captured monitor instance was ready and the MSMON_CSU_CAPTURE.VALUE field is accurate.
- 0b1 The captured monitor instance was not ready and the contents of the MSMON_CSU_CAPTURE.VALUE field might be inaccurate or otherwise not represent the actual cache storage usage.

VALUE, bits [30:0]

Captured cache storage usage measurement if MSMON_CSU_CAPTURE.NRDY is 0. Invalid if MSMON_CSU_CAPTURE.NRDY is 1.

VALUE is the captured cache storage usage measurement in bytes meeting the criteria set in [MSMON_CFG_CSU_FLT](#) and [MSMON_CFG_CSU_CTL](#) for the monitor instance selected by [MSMON_CFG_MON_SEL](#).

Accessing the MSMON_CSU_CAPTURE:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_CSU_CAPTURE_s must only be accessible from the Secure MPAM feature page.
- MSMON_CSU_CAPTURE_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_CSU_CAPTURE_rt must only be accessible from the Root MPAM feature page.
- MSMON_CSU_CAPTURE_rl must only be accessible from the Realm MPAM feature page.

MSMON_CSU_CAPTURE_s, MSMON_CSU_CAPTURE_ns, MSMON_CSU_CAPTURE_rt, and MSMON_CSU_CAPTURE_rl must be separate registers:

- The Secure instance (MSMON_CSU_CAPTURE_s) accesses the captured cache storage usage monitor used for Secure PARTIDs.
- The Non-secure instance (MSMON_CSU_CAPTURE_ns) accesses the captured cache storage usage monitor used for Non-secure PARTIDs.
- The Root instance (MSMON_CSU_CAPTURE_rt) accesses the captured cache storage usage monitor used for Root PARTIDs.
- The Realm instance (MSMON_CSU_CAPTURE_rl) accesses the captured cache storage usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to MSMON_CSU_CAPTURE access the monitor instance for the cache resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the cache storage usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When RIS is not implemented, reads and writes to MSMON_CSU_CAPTURE access the monitor instance for the cache storage usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

MSMON_CSU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0848	MSMON_CSU_CAPTURE_s

Accesses to this interface are RW.

MSMON_CSU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0848	MSMON_CSU_CAPTURE_ns

Accesses to this interface are RW.

MSMON_CSU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0848	MSMON_CSU_CAPTURE_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_CSU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0848	MSMON_CSU_CAPTURE_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.9 MSMON_CSU_OFSR, MPAM CSU Monitor Overflow Status Register

The MSMON_CSU_OFSR characteristics are:

Purpose

MSMON_CSU_OFSR is a 32-bit read-only register that shows bitmap of CSU monitor instance overflow status for a contiguous group of 32 monitor instances.

MSMON_CSU_OFSR_s gives a bitmap of pending CSU overflow status for 32 Secure CSU monitor instances. MSMON_CSU_OFSR_ns gives a bitmap of pending CSU overflow status for 32 Non-secure CSU monitor instances. MSMON_CSU_OFSR_rt gives a bitmap of pending CSU overflow status for 32 Root CSU monitor instances. MSMON_CSU_OFSR_rl gives a bitmap of pending CSU overflow status for 32 Realm CSU monitor instances.

Configurations

The power domain of MSMON_CSU_OFSR is IMPLEMENTATION DEFINED.

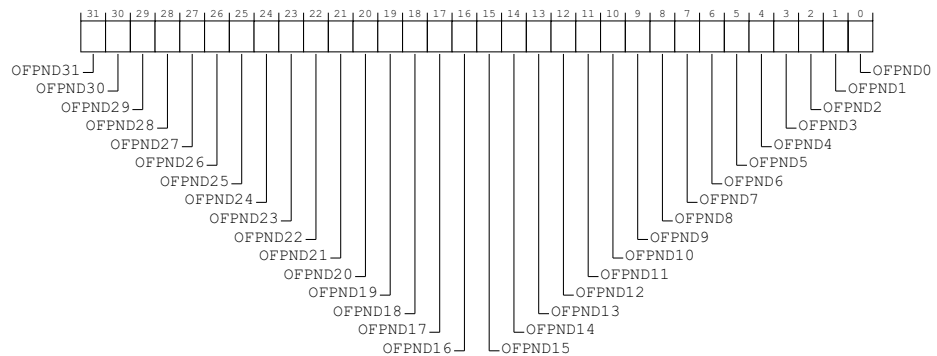
This register is present only when MPAMF_CSUMON_IDR.HAS_OFSR == 1. Otherwise, direct accesses to MSMON_CSU_OFSR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_CSU_OFSR is a 32-bit register.

Field descriptions



OFPND<i>, bit [i], for i = 31 to 0

Overflow status bitmap for CSU monitor instances. The RIS and the contiguous range of CSU monitor instances are set in MSMON_CFG_MON_SEL. i of 0 corresponds to the CSU monitor instance MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0.

0b0 CSU monitor instance (MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0 + i) does not have a pending overflow.

0b1 CSU monitor instance (MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0 + i) has a pending overflow.

After reading MSMON_OFLOW_SR to determine that a CSU monitor instance has a pending overflow and which RIS values have pending overflows, an interrupt service routine could poll groups of 32 monitor instances in a RIS for pending monitors by reading this bitmap and incrementing MSMON_CFG_MON_SEL.MON_SEL by 32.

Accessing the MSMON_CSU_OFSR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_CSU_OFSR_s must only be accessible from the Secure MPAM feature page.
- MSMON_CSU_OFSR_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_CSU_OFSR_rt must only be accessible from the Root MPAM feature page.
- MSMON_CSU_OFSR_rl must only be accessible from the Realm MPAM feature page.

MSMON_CSU_OFSR_s, MSMON_CSU_OFSR_ns, MSMON_CSU_OFSR_rt, and MSMON_CSU_OFSR_rl must be separate registers:

- The Secure instance (MSMON_CSU_OFSR_s) accesses the CSU monitor overflow status bitmap used for Secure PARTIDs.
- The Non-secure instance (MSMON_CSU_OFSR_ns) accesses the CSU monitor overflow status bitmap used for Non-secure PARTIDs.
- The Root instance (MSMON_CSU_OFSR_rt) accesses the CSU monitor overflow status bitmap used for Root PARTIDs.
- The Realm instance (MSMON_CSU_OFSR_rl) accesses the CSU monitor overflow status bitmap used for Realm PARTIDs.

MSMON_CSU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0858	MSMON_CSU_OFSR_s

Accesses to this interface are RO.

MSMON_CSU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0858	MSMON_CSU_OFSR_ns

Accesses to this interface are RO.

MSMON_CSU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0858	MSMON_CSU_OFSR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MSMON_CSU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0858	MSMON_CSU_OFSR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.5.10 MSMON_MBWU, MPAM Memory Bandwidth Usage Monitor Register

The MSMON_MBWU characteristics are:

Purpose

Accesses the monitor instance selected by `MSMON_CFG_MON_SEL`.

`MSMON_MBWU_s` is the Secure memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_s`. `MSMON_MBWU_ns` is the Non-secure memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_ns`. `MSMON_MBWU_rt` is the Root memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_rt`. `MSMON_MBWU_rl` is the Realm memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_rl`.

If `MPAMF_IDR.HAS_RIS` is 1, the monitor instance register accessed is for the resource instance currently selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance of that resource instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

Configurations

The power domain of `MSMON_MBWU` is IMPLEMENTATION DEFINED.

This register is present only when `FEAT_MPAM` is implemented, `MPAMF_IDR.HAS_MSMON == 1` and `MPAMF_MSMON_IDR.MSMON_MBWU == 1`. Otherwise, direct accesses to `MSMON_MBWU` are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

`MSMON_MBWU` is a 32-bit register.

Field descriptions



NRDY, bit [31]

Not Ready. Indicates whether the monitor has possibly inaccurate data.

0b0 The monitor instance is ready and the `MSMON_MBWU.VALUE` field is accurate.

0b1 The monitor instance is not ready and the contents of the `MSMON_MBWU.VALUE` field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

VALUE, bits [30:0]

Memory bandwidth usage counter value if `MSMON_MBWU.NRDY` is 0. Invalid if `MSMON_MBWU.NRDY` is 1.

`VALUE` is the scaled count of bytes transferred since the monitor was last reset that met the criteria set in `MSMON_CFG_MBWU_FLT` and `MSMON_CFG_MBWU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

If `MSMON_CFG_MBWU_CTL.SCLEN` enables scaling, the count in `VALUE` is the number of bytes shifted right by `MPAMF_MBWUMON_IDR.SCALE` bit positions and rounded.

Accessing the MSMON_MBWU:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_MBWU_s must only be accessible from the Secure MPAM feature page.
- MSMON_MBWU_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_MBWU_rt must only be accessible from the Root MPAM feature page.
- MSMON_MBWU_rl must only be accessible from the Realm MPAM feature page.

MSMON_MBWU_s, MSMON_MBWU_ns, MSMON_MBWU_rt, and MSMON_MBWU_rl must be separate registers:

- The Secure instance (MSMON_MBWU_s) accesses the memory bandwidth usage monitor used for Secure PARTIDs.
- The Non-secure instance (MSMON_MBWU_ns) accesses the memory bandwidth usage monitor used for Non-secure PARTIDs.
- The Root instance (MSMON_MBWU_rt) accesses the memory bandwidth usage monitor used for Root PARTIDs.
- The Realm instance (MSMON_MBWU_rl) accesses the memory bandwidth usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to MSMON_MBWU access the memory bandwidth usage monitor instance for the resource instance selected by [MSMON_CFG_MON_SEL.RIS](#) and the memory bandwidth usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

When RIS is not implemented, reads and writes to MSMON_MBWU access the memory bandwidth usage monitor instance for the memory bandwidth usage monitor instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

MSMON_MBWU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0860	MSMON_MBWU_s

Accesses to this interface are RW.

MSMON_MBWU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0860	MSMON_MBWU_ns

Accesses to this interface are RW.

MSMON_MBWU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0860	MSMON_MBWU_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_MBWU can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0860	MSMON_MBWU_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.11 MSMON_MBWU_CAPTURE, MPAM Memory Bandwidth Usage Monitor Capture Register

The MSMON_MBWU_CAPTURE characteristics are:

Purpose

Accesses the captured MSMON_MBWU monitor instance selected by [MSMON_CFG_MON_SEL](#).

MSMON_MBWU_CAPTURE_s is the Secure memory bandwidth usage monitor capture instance selected by the Secure instance of [MSMON_CFG_MON_SEL](#). MSMON_MBWU_CAPTURE_ns is the Non-secure memory bandwidth usage monitor capture instance selected by the Non-secure instance of [MSMON_CFG_MON_SEL](#). MSMON_MBWU_CAPTURE_rt is the Root memory bandwidth usage monitor capture instance selected by the Root instance of [MSMON_CFG_MON_SEL](#). MSMON_MBWU_CAPTURE_rl is the Realm memory bandwidth usage monitor capture instance selected by the Realm instance of [MSMON_CFG_MON_SEL](#).

If [MPAMF_IDR.HAS_RIS](#) is 1, the monitor instance capture register accessed is for the resource instance currently selected by [MSMON_CFG_MON_SEL.RIS](#) and the monitor instance of that resource instance selected by [MSMON_CFG_MON_SEL.MON_SEL](#).

Configurations

The power domain of MSMON_MBWU_CAPTURE is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, [MPAMF_IDR.HAS_MSMON](#) == 1, [MPAMF_MSMON_IDR.MSMON_MBWU](#) == 1 and [MPAMF_MBWUMON_IDR.HAS_CAPTURE](#) == 1. Otherwise, direct accesses to MSMON_MBWU_CAPTURE are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_MBWU_CAPTURE is a 32-bit register.

Field descriptions



NRDY, bit [31]

Not Ready. The captured NRDY bit from the corresponding instance of [MSMON_MBWU](#). This bit indicates whether the captured monitor value has possibly inaccurate data.

- 0b0 The captured monitor instance was ready and the [MSMON_MBWU_CAPTURE.VALUE](#) field is accurate.
- 0b1 The captured monitor instance was not ready and the contents of the [MSMON_MBWU_CAPTURE.VALUE](#) field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

VALUE, bits [30:0]

Captured memory bandwidth usage counter value if [MSMON_MBWU_CAPTURE.NRDY](#) is 0. Invalid if [MSMON_MBWU_CAPTURE.NRDY](#) is 1.

VALUE is the captured VALUE field from the corresponding instance of [MSMON_MBWU](#), the count of bytes transferred since the monitor was last reset that meet the criteria set in [MSMON_CFG_MBWU_FLT](#) and [MSMON_CFG_MBWU_CTL](#) for the monitor instance selected by [MSMON_CFG_MON_SEL](#).

VALUE captures the [MSMON_MBWU.VALUE](#) and preserves any scaling that had been performed on the VALUE field in that register.

Accessing the MSMON_MBWU_CAPTURE:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_MBWU_CAPTURE_s must only be accessible from the Secure MPAM feature page.
- MSMON_MBWU_CAPTURE_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_MBWU_CAPTURE_rt must only be accessible from the Root MPAM feature page.
- MSMON_MBWU_CAPTURE_rl must only be accessible from the Realm MPAM feature page.

MSMON_MBWU_CAPTURE_s, MSMON_MBWU_CAPTURE_ns, MSMON_MBWU_CAPTURE_rt, and MSMON_MBWU_CAPTURE_rl must be separate registers:

- The Secure instance (MSMON_MBWU_CAPTURE_s) accesses the captured memory bandwidth usage monitor used for Secure PARTIDs.
- The Non-secure instance (MSMON_MBWU_CAPTURE_ns) accesses the captured memory bandwidth usage monitor used for Non-secure PARTIDs.
- The Root instance (MSMON_MBWU_CAPTURE_rt) accesses the captured memory bandwidth usage monitor used for Root PARTIDs.
- The Realm instance (MSMON_MBWU_CAPTURE_rl) accesses the captured memory bandwidth usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to MSMON_MBWU_CAPTURE access the monitor instance for the bandwidth resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When RIS is not implemented, reads and writes to MSMON_MBWU_CAPTURE access the monitor instance for the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

MSMON_MBWU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0868	MSMON_MBWU_CAPTURE_s

Accesses to this interface are RW.

MSMON_MBWU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0868	MSMON_MBWU_CAPTURE_ns

Accesses to this interface are RW.

MSMON_MBWU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0868	MSMON_MBWU_CAPTURE_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_MBWU_CAPTURE can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_r1	0x0868	MSMON_MBWU_CAPTURE_r1

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.12 MSMON_MBWU_L, MPAM Long Memory Bandwidth Usage Monitor Register

The MSMON_MBWU_L characteristics are:

Purpose

Accesses the monitor instance selected by `MSMON_CFG_MON_SEL`.

`MSMON_MBWU_L_s` is the Secure long memory bandwidth usage monitor instance selected by the Secure instance of `MSMON_CFG_MON_SEL`. `MSMON_MBWU_L_ns` is the Non-secure long memory bandwidth usage monitor instance selected by the Non-secure instance of `MSMON_CFG_MON_SEL`. `MSMON_MBWU_L_rt` is the Root long memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_rt`. `MSMON_MBWU_L_rl` is the Realm long memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL_rl`.

If `MPAMF_IDR.HAS_RIS` is 1, the monitor instance long monitor register accessed is for the resource instance currently selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance of that resource instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

Configurations

The power domain of `MSMON_MBWU_L` is IMPLEMENTATION DEFINED.

This register is present only when `FEAT_MPAM` is implemented, `MPAMF_IDR.HAS_MSMON == 1`, `MPAMF_MSMON_IDR.MSMON_MBWU == 1` and `MPAMF_MBWUMON_IDR.HAS_LONG == 1`. Otherwise, direct accesses to `MSMON_MBWU_L` are `RES0`.


The power and reset domain of each MSC component is specific to that component.

Attributes

`MSMON_MBWU_L` is a 64-bit register.

Field descriptions

When `MPAMF_MBWUMON_IDR.LWD == 0`:



NRDY, bit [63]

Not Ready. Indicates whether the monitor instance has possibly inaccurate data.

- 0b0 The monitor instance is ready and the `MSMON_MBWU_L.VALUE` field is accurate.
- 0b1 The monitor instance is not ready and the contents of the `MSMON_MBWU_L.VALUE` field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

Bits [62:44]


Reserved, `RES0`.

VALUE, bits [43:0]

Long (44-bit) memory bandwidth usage counter value if `MSMON_MBWU_L.NRDY` is 0. Invalid if `MSMON_MBWU_L.NRDY` is 1.

`VALUE` is the long count of bytes transferred since the monitor was last reset that met the criteria set in `MSMON_CFG_MBWU_FLT` and `MSMON_CFG_MBWU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

When `MPAMF_MBWUMON_IDR.LWD == 1`:



NRDY, bit [63]

Not Ready. Indicates whether the monitor instance has possibly inaccurate data.

- 0b0 The monitor instance is ready and the `MSMON_MBWU_L.VALUE` field is accurate.
- 0b1 The monitor instance is not ready and the contents of the `MSMON_MBWU_L.VALUE` field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

VALUE, bits [62:0]

Long (63-bit) memory bandwidth usage counter value if `MSMON_MBWU_L.NRDY` is 0. Invalid if `MSMON_MBWU_L.NRDY` is 1.

VALUE is the long count of bytes transferred since the monitor instance was last reset that met the criteria set in `MSMON_CFG_MBWU_FLT` and `MSMON_CFG_MBWU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

Accessing the `MSMON_MBWU_L`:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MSMON_MBWU_L_s` must only be accessible from the Secure MPAM feature page.
- `MSMON_MBWU_L_ns` must only be accessible from the Non-secure MPAM feature page.
- `MSMON_MBWU_L_rt` must only be accessible from the Root MPAM feature page.
- `MSMON_MBWU_L_rl` must only be accessible from the Realm MPAM feature page.

`MSMON_MBWU_L_s`, `MSMON_MBWU_L_ns`, `MSMON_MBWU_L_rt`, and `MSMON_MBWU_L_rl` must be separate registers:

- The Secure instance (`MSMON_MBWU_L_s`) accesses the long memory bandwidth usage monitor used for Secure PARTIDs.
- The Non-secure instance (`MSMON_MBWU_L_ns`) accesses the long memory bandwidth usage monitor used for Non-secure PARTIDs.
- The Root instance (`MSMON_MBWU_L_rt`) accesses the long memory bandwidth usage monitor used for Root PARTIDs.
- The Realm instance (`MSMON_MBWU_L_rl`) accesses the long memory bandwidth usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to `MSMON_MBWU_L` access the long memory bandwidth usage monitor instance for the bandwidth resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When RIS is not implemented, reads and writes to `MSMON_MBWU_L` access the long memory bandwidth usage monitor instance for the monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

MSMON_MBWU_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0880	MSMON_MBWU_s

Accesses to this interface are RW.

MSMON_MBWU_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0880	MSMON_MBWU_ns

Accesses to this interface are RW.

MSMON_MBWU_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0880	MSMON_MBWU_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_MBWU_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0880	MSMON_MBWU_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.13 MSMON_MBWU_L_CAPTURE, MPAM Long Memory Bandwidth Usage Monitor Capture Register

The MSMON_MBWU_L_CAPTURE characteristics are:

Purpose

Accesses the captured MSMON_MBWU_L monitor instance selected by MSMON_CFG_MON_SEL.

MSMON_MBWU_L_CAPTURE_s is the Secure long memory bandwidth usage monitor capture instance selected by the Secure instance of MSMON_CFG_MON_SEL.

MSMON_MBWU_L_CAPTURE_ns is the Non-secure long memory bandwidth usage monitor capture instance selected by the Non-secure instance of MSMON_CFG_MON_SEL.

MSMON_MBWU_L_CAPTURE_rt is the Root long memory bandwidth usage monitor capture instance selected by the Root instance of MSMON_CFG_MON_SEL.

MSMON_MBWU_L_CAPTURE_rl is the Realm long memory bandwidth usage monitor capture instance selected by the Realm instance of MSMON_CFG_MON_SEL.

If MPAMF_IDR.HAS_RIS is 1, the monitor instance long capture register accessed is for the resource instance currently selected by MSMON_CFG_MON_SEL.RIS and the monitor instance of that resource instance selected by MSMON_CFG_MON_SEL.MON_SEL.

Configurations

The power domain of MSMON_MBWU_L_CAPTURE is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented, MPAMF_IDR.HAS_MSMON == 1, MPAMF_MSMON_IDR.MSMON_MBWU == 1, MPAMF_MBWUMON_IDR.HAS_CAPTURE == 1 and MPAMF_MBWUMON_IDR.HAS_LONG == 1. Otherwise, direct accesses to MSMON_MBWU_L_CAPTURE are RES0.


The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_MBWU_L_CAPTURE is a 64-bit register.

Field descriptions

When MPAMF_MBWUMON_IDR.LWD == 0:



NRDY, bit [63]

Not Ready. Indicates whether the monitor has possibly inaccurate data.

0b0 The captured monitor instance was ready and the MSMON_MBWU_L_CAPTURE.VALUE field is accurate.

0b1 The captured monitor instance was not ready and the contents of the MSMON_MBWU_L_CAPTURE.VALUE field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

Bits [62:44]


Reserved, RES0.

VALUE, bits [43:0]

Captured long memory bandwidth usage counter value if MSMON_MBWU_L_CAPTURE.NRDY is 0. Invalid if MSMON_MBWU_L_CAPTURE.NRDY is 1.

VALUE is the captured 44-bit count of bytes transferred since the monitor instance was last reset that met the criteria set in `MSMON_CFG_MBWU_FLT` and `MSMON_CFG_MBWU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

When `MPAMF_MBWUMON_IDR.LWD == 1`:



NRDY, bit [63]

Not Ready. Indicates whether the monitor has possibly inaccurate data.

- 0b0 The captured monitor instance was ready and the `MSMON_MBWU_L_CAPTURE.VALUE` field is accurate.
- 0b1 The captured monitor instance was not ready and the contents of the `MSMON_MBWU_L_CAPTURE.VALUE` field might be inaccurate or otherwise not represent the actual memory bandwidth usage.

VALUE, bits [62:0]

The captured long memory bandwidth usage counter value if `MSMON_MBWU_L_CAPTURE.NRDY` is 0. Invalid if `MSMON_MBWU_L_CAPTURE.NRDY` is 1.

VALUE is the captured 63-bit count of bytes transferred since the monitor instance was last reset that met the criteria set in `MSMON_CFG_MBWU_FLT` and `MSMON_CFG_MBWU_CTL` for the monitor instance selected by `MSMON_CFG_MON_SEL`.

Accessing the `MSMON_MBWU_L_CAPTURE`:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- `MSMON_MBWU_L_CAPTURE_s` must only be accessible from the Secure MPAM feature page.
- `MSMON_MBWU_L_CAPTURE_ns` must only be accessible from the Non-secure MPAM feature page.
- `MSMON_MBWU_L_CAPTURE_rt` must only be accessible from the Root MPAM feature page.
- `MSMON_MBWU_L_CAPTURE_rl` must only be accessible from the Realm MPAM feature page.

`MSMON_MBWU_L_CAPTURE_s`, `MSMON_MBWU_L_CAPTURE_ns`, `MSMON_MBWU_L_CAPTURE_rt`, and `MSMON_MBWU_L_CAPTURE_rl` must be separate registers:

- The Secure instance (`MSMON_MBWU_L_CAPTURE_s`) accesses the captured long memory bandwidth usage monitor used for Secure PARTIDs.
- The Non-secure instance (`MSMON_MBWU_L_CAPTURE_ns`) accesses the captured long memory bandwidth usage monitor used for Non-secure PARTIDs.
- The Root instance (`MSMON_MBWU_L_CAPTURE_rt`) accesses the captured long memory bandwidth usage monitor used for Root PARTIDs.
- The Realm instance (`MSMON_MBWU_L_CAPTURE_rl`) accesses the captured long memory bandwidth usage monitor used for Realm PARTIDs.

When RIS is implemented, reads and writes to `MSMON_MBWU_L_CAPTURE` access the monitor instance for the bandwidth resource instance selected by `MSMON_CFG_MON_SEL.RIS` and the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

When RIS is not implemented, reads and writes to `MSMON_MBWU_L_CAPTURE` access the monitor instance for the memory bandwidth usage monitor instance selected by `MSMON_CFG_MON_SEL.MON_SEL`.

`MSMON_MBWU_L_CAPTURE` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0890	MSMON_MBWU_CAPTURE_s

Accesses to this interface are RW.

`MSMON_MBWU_L_CAPTURE` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0890	MSMON_MBWU_CAPTURE_ns

Accesses to this interface are RW.

`MSMON_MBWU_L_CAPTURE` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0890	MSMON_MBWU_CAPTURE_rt

When `FEAT_RME` is implemented, accesses to this interface are RW.

`MSMON_MBWU_L_CAPTURE` can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0890	MSMON_MBWU_CAPTURE_rl

When `FEAT_RME` is implemented, accesses to this interface are RW.

9.5.14 MSMON_MBWU_OFSR, MPAM MBWU Monitor Overflow Status Register

The MSMON_MBWU_OFSR characteristics are:

Purpose

MSMON_MBWU_OFSR is a 32-bit read-only register that shows bitmap of MBWU monitor instance overflow status for a contiguous group of 32 monitor instances.

MSMON_MBWU_OFSR_s gives a bitmap of pending MBWU overflow status for 32 Secure MBWU monitor instances. MSMON_MBWU_OFSR_ns gives a bitmap of pending MBWU overflow status for 32 Non-secure MBWU monitor instances. MSMON_MBWU_OFSR_rt gives a bitmap of pending MBWU overflow status for 32 Root MBWU monitor instances. MSMON_MBWU_OFSR_rl gives a bitmap of pending MBWU overflow status for 32 Realm MBWU monitor instances.

Configurations

The power domain of MSMON_MBWU_OFSR is IMPLEMENTATION DEFINED.


This register is present only when MPAMF_MBWUMON_IDR.HAS_OFSR == 1. Otherwise, direct accesses to MSMON_MBWU_OFSR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_MBWU_OFSR is a 32-bit register.

Field descriptions



OFPND<i>, bit [i], for i = 31 to 0

Overflow status bitmap for MBWU monitor instances. The RIS and the contiguous range of MBWU monitor instances are set in `MSMON_CFG_MON_SEL`. `i` of 0 corresponds to the MBWU monitor instance `MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0`.

- 0b0 MBWU monitor instance (`MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0 + i`) does not have a pending overflow.
- 0b1 MBWU monitor instance (`MSMON_CFG_MON_SEL.MON_SEL & 0xFFE0 + i`) has a pending overflow.

After reading `MSMON_OFLOW_SR` to determine that an MBWU monitor instance has a pending overflow and which RIS values have pending overflows, an interrupt service routine could poll groups of 32 monitor instances in a RIS for pending monitors by reading this bitmap and incrementing `MSMON_CFG_MON_SEL.MON_SEL` by 32.

A pending overflow may be in either the `MSMON_CFG_MBWU_CTL.OFLOW_STATUS` or `MSMON_CFG_MBWU_CTL.OFLOW_STATUS_L` field.

Accessing the MSMON_MBWU_OFSR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_MBWU_OFSR_s must only be accessible from the Secure MPAM feature page.
- MSMON_MBWU_OFSR_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_MBWU_OFSR_rt must only be accessible from the Root MPAM feature page.
- MSMON_MBWU_OFSR_rl must only be accessible from the Realm MPAM feature page.

MSMON_MBWU_OFSR_s, MSMON_MBWU_OFSR_ns, MSMON_MBWU_OFSR_rt, and MSMON_MBWU_OFSR_rl must be separate registers:

- The Secure instance (MSMON_MBWU_OFSR_s) accesses the MBWU monitor overflow status bitmap used for Secure PARTIDs.
- The Non-secure instance (MSMON_MBWU_OFSR_ns) accesses the MBWU monitor overflow status bitmap used for Non-secure PARTIDs.
- The Root instance (MSMON_MBWU_OFSR_rt) accesses the MBWU monitor overflow status bitmap used for Root PARTIDs.
- The Realm instance (MSMON_MBWU_OFSR_rl) accesses the MBWU monitor overflow status bitmap used for Realm PARTIDs.

MSMON_MBWU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x0898	MSMON_MBWU_OFSR_s

Accesses to this interface are RO.

MSMON_MBWU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x0898	MSMON_MBWU_OFSR_ns

Accesses to this interface are RO.

MSMON_MBWU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x0898	MSMON_MBWU_OFSR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MSMON_MBWU_OFSR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x0898	MSMON_MBWU_OFSR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.5.15 MSMON_OFLOW_MSI_ADDR_H, MPAM Monitor Overflow MSI Write High-part Address Register

The MSMON_OFLOW_MSI_ADDR_H characteristics are:

Purpose

MSMON_OFLOW_MSI_ADDR_H is a 32-bit read/write register for the high part of the MPAM monitor overflow MSI address.

MSMON_OFLOW_MSI_ADDR_H_s is the high part of the MSI write address for monitor overflow interrupts from Secure monitor instances. MSMON_OFLOW_MSI_ADDR_H_ns is the high part of the MSI write address for monitor overflow interrupts from Non-secure monitor instances. MSMON_OFLOW_MSI_ADDR_H_rt is the high part of the MSI write address for monitor overflow interrupts from Root monitor instances. MSMON_OFLOW_MSI_ADDR_H_rl is the high part of the MSI write address for monitor overflow interrupts from Realm monitor instances.

Configurations

The power domain of MSMON_OFLOW_MSI_ADDR_H is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAMv1p1 is implemented and MPAMF_MSMON_IDR.HAS_OFLW_MSI == 1. Otherwise, direct accesses to MSMON_OFLOW_MSI_ADDR_H are RES0.

MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, and MSMON_OFLOW_MSI_MPAM must all be implemented to support MSI writes for monitor overflow interrupts.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_MSI_ADDR_H is a 32-bit register.

Field descriptions



Bits [31:20]

Reserved, RES0.

MSI_ADDR_H, bits [19:0]

MSI write address bits[51:32].

Accessing the MSMON_OFLOW_MSI_ADDR_H:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_MSI_ADDR_H_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_H_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_H_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_H_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLW_MSI_ADDR_H_s, MSMON_OFLW_MSI_ADDR_H_ns, MSMON_OFLW_MSI_ADDR_H_rt, and MSMON_OFLW_MSI_ADDR_H_rl must be separate registers:

- The Secure instance (MSMON_OFLW_MSI_ADDR_H_s) accesses the high part of the monitor overflow MSI write address of Secure monitors.
- The Non-secure instance (MSMON_OFLW_MSI_ADDR_H_ns) accesses the high part of the monitor overflow MSI write address of Non-secure monitors.
- The Root instance (MSMON_OFLW_MSI_ADDR_H_rt) accesses the high part of the monitor overflow MSI write address of Root monitors.
- The Realm instance (MSMON_OFLW_MSI_ADDR_H_rl) accesses the high part of the monitor overflow MSI write address of Realm monitors.

MSMON_OFLOW_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08E4	MSMON_OFLW_MSI_ADDR_H_s

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08E4	MSMON_OFLW_MSI_ADDR_H_ns

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08E4	MSMON_OFLW_MSI_ADDR_H_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08E4	MSMON_OFLW_MSI_ADDR_H_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.16 MSMON_OFLOW_MSI_ADDR_L, MPAM Monitor Overflow MSI Low-part Address Register

The MSMON_OFLOW_MSI_ADDR_L characteristics are:

Purpose

MSMON_OFLOW_MSI_ADDR_L is a 32-bit read/write register for the low part of the MPAM monitor MSI address.

MSMON_OFLOW_MSI_ADDR_L_s is the low part of the MSI write address for overflow interrupts from Secure monitor instances. MSMON_OFLOW_MSI_ADDR_L_ns is the low part of the MSI write address for overflow interrupts from Non-secure monitor instances.

MSMON_OFLOW_MSI_ADDR_L_rt is the low part of the MSI write address for overflow interrupts from Root monitor instances. MSMON_OFLOW_MSI_ADDR_L_rl is the low part of the MSI write address for overflow interrupts from Realm monitor instances.

Configurations

The power domain of MSMON_OFLOW_MSI_ADDR_L is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAMv1p1 is implemented and MPAMF_MSMON_IDR.HAS_OFLW_MSI == 1. Otherwise, direct accesses to MSMON_OFLOW_MSI_ADDR_L are RES0.

MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, and MSMON_OFLOW_MSI_MPAM must all be implemented to support MSI writes for monitor overflow interrupts.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_MSI_ADDR_L is a 32-bit register.

Field descriptions



MSI_ADDR_L, bits [31:2]

MSI write address bits[31:2].

Bits [1:0]

Reads as 0b00.

Access to this field is RO.

Accessing the MSMON_OFLOW_MSI_ADDR_L:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_MSI_ADDR_L_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_L_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_L_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_MSI_ADDR_L_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLOW_MSI_ADDR_L_s, MSMON_OFLOW_MSI_ADDR_L_ns, MSMON_OFLOW_MSI_ADDR_L_rt, and MSMON_OFLOW_MSI_ADDR_L_rl must be separate registers:

- The Secure instance (MSMON_OFLOW_MSI_ADDR_L_s) accesses the low part of the overflow MSI write address used for Secure PARTIDs.
- The Non-secure instance (MSMON_OFLOW_MSI_ADDR_L_ns) accesses the low part of the overflow MSI write address used for Non-secure PARTIDs.
- The Root instance (MSMON_OFLOW_MSI_ADDR_L_rt) accesses the low part of the overflow MSI write address used for Root PARTIDs.
- The Realm instance (MSMON_OFLOW_MSI_ADDR_L_rl) accesses the low part of the overflow MSI write address used for Realm PARTIDs.

MSMON_OFLOW_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08E0	MSMON_OFLOW_MSI_ADDR_L_s

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08E0	MSMON_OFLOW_MSI_ADDR_L_ns

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08E0	MSMON_OFLOW_MSI_ADDR_L_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_OFLOW_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08E0	MSMON_OFLOW_MSI_ADDR_L_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.17 MSMON_OFLOW_MSI_ATTR, MPAM Monitor Overflow MSI Write Attributes Register

The MSMON_OFLOW_MSI_ATTR characteristics are:

Purpose

MSMON_OFLOW_MSI_ATTR is a 32-bit read/write register that controls MPAM monitor overflow MSI write attributes for MPAM monitor overflows in this MSC.

MSMON_OFLOW_MSI_ATTR_s controls Secure MPAM monitor overflow MSI writes.
 MSMON_OFLOW_MSI_ATTR_ns controls Non-secure MPAM monitor overflow MSI writes.
 MSMON_OFLOW_MSI_ATTR_rt controls Root MPAM monitor overflow MSI writes.
 MSMON_OFLOW_MSI_ATTR_rl controls Realm MPAM monitor overflow MSI writes.

Configurations

The power domain of MSMON_OFLOW_MSI_ATTR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAMv1p1 is implemented and MPAMF_MSMON_IDR.HAS_OFLW_MSI == 1. Otherwise, direct accesses to MSMON_OFLOW_MSI_ATTR are RES0.

MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, and MSMON_OFLOW_MSI_MPAM must all be implemented to support MSI writes for monitor overflow interrupts.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_MSI_ATTR is a 32-bit register.

Field descriptions



Bits [31:30]

Reserved, RES0.

MSI_SH, bits [29:28]

Sharability attribute of MSI writes.

- 0b00 Non-shareable.
- 0b01 Reserved, CONSTRAINED UNPREDICTABLE.
- 0b10 Outer Shareable.
- 0b11 Inner Shareable.

When MSMON_OFLOW_MSI_ATTR.MSI_MEMATTR specifies a Device memory type, the contents of this field are IGNORED and Shareability is effectively Outer Shareable.

MSI_MEMATTR, bits [27:24]

Memory attributes of MSI writes.

Note: This encoding matches the VMSAv8-64 stage 2 MemAttr[3:0] field as described in the Arm ARM, except that the following encodings are Reserved (not UNPREDICTABLE) and behave as Device-nGnRnE: 0b0100, 0b1000, and 0b1100.

- 0b0000 Device-nGnRnE.
- 0b0001 Device-nGnRE.
- 0b0010 Device-nGRE.

0b0011	Device-GRE.
0b0100	Reserved. Behave as Device-nGnRnE, 0b0000.
0b0101	Normal Inner Non-cacheable, Outer Non-cacheable.
0b0110	Normal Inner Write-Through Cacheable, Outer Non-cacheable.
0b0111	Normal Inner Write-Back Cacheable, Outer Non-cacheable.
0b1000	Reserved. Behave as Device-nGnRnE, 0b0000.
0b1001	Normal Inner Non-Cachable, Outer Write-Through Cacheable.
0b1010	Normal Inner Write-Through Cacheable, Outer Write-Through Cachable.
0b1011	Normal Inner Write-Back Cacheable, Outer Write-Through Cachable.
0b1100	Reserved. Behave as Device-nGnRnE, 0b0000.
0b1101	Normal Inner Non-cacheable, Outer Write-Back Cacheable.
0b1110	Normal Inner Write-Through Cacheable, Outer Write-Back Cacheable.
0b1111	Normal Inner Write-Back Cacheable, Outer Write-Back Cacheable.

When this field specifies a Device memory type, the contents of MSMON_OFLOW_MSI_ATTR.MSI_SH are IGNORED and Shareability is effectively Outer Shareable.

Device types may be implemented as any Device type with more n characters. For example, if this field is set to 0b0010, an implementation may treat the MSI write as the specified type, Device-nGRE, or as Device-nGnRE or as Device-nGnRnE.

Reserved encodings 0b0100, 0b1000, and 0b1100 must be implemented to behave the same as the 0b0000 encoding.

Bits [23:1]

Reserved, RES0.

MSIEN, bit [0]

Monitor overflow MSI write enable.

0b0	MPAM monitor overflow MSI writes are not generated to signal enabled MPAM monitor overflow interrupts. When monitor overflow MSI writes are disabled, hardwired monitor overflow interrupt could be generated if hardwired monitor overflow interrupt is implemented.
0b1	MPAM monitor overflow MSI writes are generated to signal enabled MPAM monitor overflow interrupts. When monitor overflow MSI writes are enabled, hardwired monitor overflow interrupts are not generated.

This enable affects whether a hardwired overflow interrupt is generated.

The reset behavior of this field is:

- On a MSC reset, this field resets to 0.

Accessing the MSMON_OFLOW_MSI_ATTR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_MSI_ATTR_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_MSI_ATTR_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_MSI_ATTR_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_MSI_ATTR_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLOW_MSI_ATTR_s, MSMON_OFLOW_MSI_ATTR_ns, MSMON_OFLOW_MSI_ATTR_rt, and MSMON_OFLOW_MSI_ATTR_rl must be separate registers:

- The Secure instance (MSMON_OFLOW_MSI_ATTR_s) accesses the monitor overflow MSI write attributes of Secure monitors.
- The Non-secure instance (MSMON_OFLOW_MSI_ATTR_ns) accesses the monitor overflow MSI write attributes of Non-secure monitors.
- The Root instance (MSMON_OFLOW_MSI_ATTR_rt) accesses the monitor overflow MSI write attributes of Root monitors.
- The Realm instance (MSMON_OFLOW_MSI_ATTR_rl) accesses the monitor overflow MSI write attributes of Realm monitors.

MSMON_OFLOW_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08EC	MSMON_OFLOW_MSI_ATTR_s

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08EC	MSMON_OFLOW_MSI_ATTR_ns

Accesses to this interface are RW.

MSMON_OFLOW_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08EC	MSMON_OFLOW_MSI_ATTR_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_OFLOW_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08EC	MSMON_OFLOW_MSI_ATTR_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.18 MSMON_OFLOW_MSI_DATA, MPAM Monitor Overflow MSI Write Data Register

The MSMON_OFLOW_MSI_DATA characteristics are:

Purpose

MSMON_OFLOW_MSI_DATA is a 32-bit read/write register for the MPAM monitor overflow MSI data.

MSMON_OFLOW_MSI_DATA_s is the data for the MSI write for monitor overflow from Secure monitor instances. MSMON_OFLOW_MSI_DATA_ns is the data for the MSI writes for monitor overflow interrupts from Non-secure monitor instances. MSMON_OFLOW_MSI_DATA_rt is the data for the MSI write for monitor overflow from Root monitor instances.

MSMON_OFLOW_MSI_DATA_rl is the data for the MSI writes for monitor overflow interrupts from Realm monitor instances.

Configurations

The power domain of MSMON_OFLOW_MSI_DATA is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAMv1p1 is implemented and MPAMF_MSMON_IDR.HAS_OFLW_MSI == 1. Otherwise, direct accesses to MSMON_OFLOW_MSI_DATA are RES0.


MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, and MSMON_OFLOW_MSI_MPAM must all be implemented to support MSI writes for monitor overflow interrupts.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_MSI_DATA is a 32-bit register.

Field descriptions



MSI_DATA, bits [31:0]

MSI write data word.

Accessing the MSMON_OFLOW_MSI_DATA:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_MSI_DATA_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_MSI_DATA_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_MSI_DATA_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_MSI_DATA_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLOW_MSI_DATA_s, MSMON_OFLOW_MSI_DATA_ns, MSMON_OFLOW_MSI_DATA_rt, and MSMON_OFLOW_MSI_DATA_rl must be separate registers:

- The Secure instance (MSMON_OFLOW_MSI_DATA_s) accesses the monitor overflow MSI write data of Secure monitors.

- The Non-secure instance (MSMON_OFLOW_MSI_DATA_ns) accesses the monitor overflow MSI write data of Non-secure monitors.
- The Root instance (MSMON_OFLOW_MSI_DATA_rt) accesses the monitor overflow MSI write data of Root monitors.
- The Realm instance (MSMON_OFLOW_MSI_DATA_rl) accesses the monitor overflow MSI write data of Realm monitors.

MSMON_OFLOW_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08E8	MSMON_OFLOW_MSI_DATA_s

Accesses to this interface are RW.

MSMON_OFLOW_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08E8	MSMON_OFLOW_MSI_DATA_ns

Accesses to this interface are RW.

MSMON_OFLOW_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08E8	MSMON_OFLOW_MSI_DATA_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_OFLOW_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08E8	MSMON_OFLOW_MSI_DATA_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.19 MSMON_OFLOW_MSI_MPAM, MPAM Monitor Overflow MSI Write MPAM Information Register

The MSMON_OFLOW_MSI_MPAM characteristics are:

Purpose

MSMON_OFLOW_MSI_MPAM is a 32-bit read/write register that sets the MPAM information for a monitor overflow MSI write.

MSMON_OFLOW_MSI_MPAM_s controls MPAM information labeling of Secure monitor overflow MSI writes. MSMON_OFLOW_MSI_MPAM_ns controls MPAM information labeling of Non-secure monitor overflow MSI writes. MSMON_OFLOW_MSI_MPAM_rt controls MPAM information labeling of Root monitor overflow MSI writes. MSMON_OFLOW_MSI_MPAM_rl controls MPAM information labeling of Realm monitor overflow MSI writes.

Configurations

The power domain of MSMON_OFLOW_MSI_MPAM is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAMv1p1 is implemented and MPAMF_MSMON_IDR.HAS_OFLOW_MSI == 1. Otherwise, direct accesses to MSMON_OFLOW_MSI_MPAM are RES0.


MSMON_OFLOW_MSI_ADDR_L, MSMON_OFLOW_MSI_ADDR_H, MSMON_OFLOW_MSI_ATTR, MSMON_OFLOW_MSI_DATA, and MSMON_OFLOW_MSI_MPAM must all be implemented to support MSI writes for monitor overflow interrupts.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_MSI_MPAM is a 32-bit register.

Field descriptions



Bits [31:24]

Reserved, RES0.

PMG, bits [23:16]

Performance monitoring group property for an MSC monitor overflow MSI write.

The reset behavior of this field is:

- On a MSC reset, this field resets to an architecturally UNKNOWN value.

PARTID, bits [15:0]

Partition ID for an MSC monitor overflow MSI write.

The PARTID in this field is in the Secure PARTID space in the MSMON_OFLOW_MSI_MPAM_s instance and in the Non-secure PARTID space in the MSMON_OFLOW_MSI_MPAM_ns instance of this register.

The reset behavior of this field is:

- On a MSC reset, this field resets to an architecturally UNKNOWN value.

Accessing the MSMON_OFLOW_MSI_MPAM:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_MSI_MPAM_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_MSI_MPAM_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_MSI_MPAM_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_MSI_MPAM_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLOW_MSI_MPAM_s, MSMON_OFLOW_MSI_MPAM_ns, MSMON_OFLOW_MSI_MPAM_rt, and MSMON_OFLOW_MSI_MPAM_rl must be separate registers:

- The Secure instance (MSMON_OFLOW_MSI_MPAM_s) accesses the monitor overflow MSI MPAM information of Secure monitors.
- The Non-secure instance (MSMON_OFLOW_MSI_MPAM_ns) accesses the monitor overflow MSI MPAM information of Non-secure monitors.
- The Root instance (MSMON_OFLOW_MSI_MPAM_rt) accesses the monitor overflow MSI MPAM information of Root monitors.
- The Realm instance (MSMON_OFLOW_MSI_MPAM_rl) accesses the monitor overflow MSI MPAM information of Realm monitors.

MSMON_OFLOW_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08DC	MSMON_OFLOW_MSI_MPAM_s

Accesses to this interface are RW.

MSMON_OFLOW_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08DC	MSMON_OFLOW_MSI_MPAM_ns

Accesses to this interface are RW.

MSMON_OFLOW_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08DC	MSMON_OFLOW_MSI_MPAM_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MSMON_OFLOW_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08DC	MSMON_OFLOW_MSI_MPAM_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.5.20 MSMON_OFLOW_SR, MPAM Monitor Overflow Status Register

The MSMON_OFLOW_SR characteristics are:

Purpose

MSMON_OFLOW_SR is a 32-bit read-only register that shows MPAM monitor overflow status for this MSC.

MSMON_OFLOW_SR_s gives the status of overflows of Secure MPAM monitors.

MSMON_OFLOW_SR_ns gives the status of overflows of Non-secure MPAM monitors.

MSMON_OFLOW_SR_rt gives the status of overflows of Root MPAM monitors.

MSMON_OFLOW_SR_rl gives the status of overflows of Realm MPAM monitors.

Configurations

The power domain of MSMON_OFLOW_SR is IMPLEMENTATION DEFINED.


This register is present only when MPAMF_MSMON_IDR.HAS_OFLOW_SR == 1. Otherwise, direct accesses to MSMON_OFLOW_SR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MSMON_OFLOW_SR is a 32-bit register.

Field descriptions



CSU_OFLOW_PND, bit [31]

At least one cache storage usage monitor has OFLOW_STATUS of 1 in [MSMON_CFG_CSU_CTL](#).

0b0 There are no cache storage usage monitor instances where [MSMON_CFG_CSU_CTL.OFLOW_STATUS](#) is 1.

0b1 [MSMON_CFG_CSU_CTL](#) for at least one of the cache storage usage monitor instances has OFLOW_STATUS set to 1.

This field clears when [MSMON_CFG_CSU_CTL.OFLOW_STATUS](#) has been reset to 0 for all CSU monitor instances in this MSC.

MBWU_OFLOW_PND, bit [30]

At least one memory bandwidth usage monitor instance has OFLOW_STATUS or OFLOW_STATUS_L of 1 in [MSMON_CFG_MBWU_CTL](#).

0b0 There are no memory bandwidth usage monitor instances where [MSMON_CFG_MBWU_CTL.OFLOW_STATUS](#) is 1.

0b1 [MSMON_CFG_MBWU_CTL](#) for at least one of the memory bandwidth usage monitor instances has either OFLOW_STATUS or OFLOW_STATUS_L set to 1.

This field clears when [MSMON_CFG_MBWU_CTL.OFLOW_STATUS](#) and [MSMON_CFG_MBWU_CTL.OFLOW_STATUS_L](#) have been reset to 0 for all MBWU monitor instances in this MSC.

Bits [29:16]

Reserved, RES0.

RIS_PND<r>, bit [r], for r = 15 to 0

Overflow status by RIS.

0b0 RIS r has no unread overflows of any type of monitor.

0b1 RIS r has at least one unread overflow in at least one of the monitor types.

Combined with the CSU_OFLOW_PND and MBWU_OFLOW_PND flags in this register, an interrupt service routine could poll only the monitor types indicated in monitors for the resource instances flagged in this field.

Bit r is set when any monitor instance of any type in resource instance r has OFLOW_STATUS or OFLOW_STATUS_L set to 1.

Accessing the MSMON_OFLOW_SR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MSMON_OFLOW_SR_s must only be accessible from the Secure MPAM feature page.
- MSMON_OFLOW_SR_ns must only be accessible from the Non-secure MPAM feature page.
- MSMON_OFLOW_SR_rt must only be accessible from the Root MPAM feature page.
- MSMON_OFLOW_SR_rl must only be accessible from the Realm MPAM feature page.

MSMON_OFLOW_SR_s, MSMON_OFLOW_SR_ns, MSMON_OFLOW_SR_rt, and MSMON_OFLOW_SR_rl must be separate registers:

- The Secure instance (MSMON_OFLOW_SR_s) accesses the monitor overflow status summary of Secure monitors.
- The Non-secure instance (MSMON_OFLOW_SR_ns) accesses the monitor overflow status summary of Non-secure monitors.
- The Root instance (MSMON_OFLOW_SR_rt) accesses the monitor overflow status summary of Root monitors.
- The Realm instance (MSMON_OFLOW_SR_rl) accesses the monitor overflow status summary of Realm monitors.

MSMON_OFLOW_SR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x08F0	MSMON_OFLOW_SR_s

Accesses to this interface are RO.

MSMON_OFLOW_SR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x08F0	MSMON_OFLOW_SR_ns

Accesses to this interface are RO.

MSMON_OFLOW_SR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x08F0	MSMON_OFLOW_SR_rt

When FEAT_RME is implemented, accesses to this interface are RO.

MSMON_OFLOW_SR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x08F0	MSMON_OFLOW_SR_rl

When FEAT_RME is implemented, accesses to this interface are RO.

9.6 Memory-mapped control and status registers

This section lists the external control and status registers.

9.6.1 MPAMF_ECR, MPAM Error Control Register

The MPAMF_ECR characteristics are:

Purpose

MPAMF_ECR is a 32-bit read/write register that controls MPAM error interrupts for this MSC. MPAMF_ECR_s controls Secure MPAM error handling. MPAMF_ECR_ns controls Non-secure MPAM error handling. MPAMF_ECR_rt controls Root MPAM error handling. MPAMF_ECR_rl controls Realm MPAM error handling.

Configurations

The power domain of MPAMF_ECR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_ECR are RES0.


If an MSC cannot encounter any of the error conditions listed in [Errors in MSCs](#), both the MPAMF_ESR and MPAMF_ECR must be RAZ/WI.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ECR is a 32-bit register.

Field descriptions



Bits [31:1]

Reserved, RES0.

INTEN, bit [0]

Interrupt Enable.

- 0b0 MPAM error interrupts are not signaled.
- 0b1 MPAM error interrupts are signaled.

Accessing the MPAMF_ECR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ECR_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ECR_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ECR_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ECR_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ECR_s, MPAMF_ECR_ns, MPAMF_ECR_rt, and MPAMF_ECR_rl must be separate registers:

- The Secure instance (MPAMF_ECR_s) accesses the error interrupt controls used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ECR_ns) accesses the error interrupt controls used for Non-secure PARTIDs.
- The Root instance (MPAMF_ECR_rt) accesses the error interrupt controls used for Root PARTIDs.

- The Realm instance (MPAMF_ECR_rl) accesses the error interrupt controls used for Realm PARTIDs.

MPAMF_ECR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00F0	MPAMF_ECR_s

Accesses to this interface are RW.

MPAMF_ECR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00F0	MPAMF_ECR_ns

Accesses to this interface are RW.

MPAMF_ECR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00F0	MPAMF_ECR_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ECR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00F0	MPAMF_ECR_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.2 MPAMF_ERR_MSI_ADDR_H, MPAM Error MSI High-part Address Register

The MPAMF_ERR_MSI_ADDR_H characteristics are:

Purpose

MPAMF_ERR_MSI_ADDR_H is a 32-bit read/write register for the high part of the MPAM error MSI address.

MPAMF_ERR_MSI_ADDR_H_s is the high part of the MSI write address for error interrupts related to Secure PARTIDs. MPAMF_ERR_MSI_ADDR_H_ns is the high part of the MSI write address for error interrupts related to Non-secure PARTIDs. MPAMF_ERR_MSI_ADDR_H_rt is the high part of the MSI write address for error interrupts related to Root PARTIDs.

MPAMF_ERR_MSI_ADDR_H_rl is the high part of the MSI write address for error interrupts related to Realm PARTIDs.

Configurations

The power domain of MPAMF_ERR_MSI_ADDR_H is IMPLEMENTATION DEFINED.

This register is present only when MPAMF_IDR.HAS_ERR_MSI == 1. Otherwise, direct accesses to MPAMF_ERR_MSI_ADDR_H are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ERR_MSI_ADDR_H is a 32-bit register.

Field descriptions



Bits [31:20]

Reserved, RES0.

MSI_ADDR_H, bits [19:0]

MSI write address bits[51:32].

Accessing the MPAMF_ERR_MSI_ADDR_H:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ERR_MSI_ADDR_H_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ERR_MSI_ADDR_H_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ERR_MSI_ADDR_H_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ERR_MSI_ADDR_H_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ERR_MSI_ADDR_H_s, MPAMF_ERR_MSI_ADDR_H_ns, MPAMF_ERR_MSI_ADDR_H_rt, and MPAMF_ERR_MSI_ADDR_H_rl must be separate registers:

- The Secure instance (MPAMF_ERR_MSI_ADDR_H_s) accesses the high part of the memory address for MSI write to signal an MPAM error used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ERR_MSI_ADDR_H_ns) accesses the high part of the memory address for MSI write to signal an MPAM error used for Non-secure PARTIDs.

- The Root instance (MPAMF_ERR_MSI_ADDR_H_rt) accesses the high part of the memory address for MSI write to signal an MPAM error used for Root PARTIDs.
- The Realm instance (MPAMF_ERR_MSI_ADDR_H_rl) accesses the high part of the memory address for MSI write to signal an MPAM error used for Realm PARTIDs.

MPAMF_ERR_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00E4	MPAMF_ERR_MSI_ADDR_H_s

Accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00E4	MPAMF_ERR_MSI_ADDR_H_ns

Accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00E4	MPAMF_ERR_MSI_ADDR_H_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_H can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00E4	MPAMF_ERR_MSI_ADDR_H_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.3 MPAMF_ERR_MSI_ADDR_L, MPAM Error MSI Low-part Address Register

The MPAMF_ERR_MSI_ADDR_L characteristics are:

Purpose

MPAMF_ERR_MSI_ADDR_L is a 32-bit read/write register for the low part of the MPAM error MSI address.

MPAMF_ERR_MSI_ADDR_L_s is the low part of the MSI write address for error interrupts related to Secure PARTIDs. MPAMF_ERR_MSI_ADDR_L_ns is the low part of the MSI write address for error interrupts related to Non-secure PARTIDs. MPAMF_ERR_MSI_ADDR_L_rt is the low part of the MSI write address for error interrupts related to Root PARTIDs.

MPAMF_ERR_MSI_ADDR_L_rl is the low part of the MSI write address for error interrupts related to Realm PARTIDs.

Configurations

The power domain of MPAMF_ERR_MSI_ADDR_L is IMPLEMENTATION DEFINED.

This register is present only when MPAMF_IDR.HAS_ERR_MSI == 1. Otherwise, direct accesses to MPAMF_ERR_MSI_ADDR_L are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ERR_MSI_ADDR_L is a 32-bit register.

Field descriptions



MSI_ADDR_L, bits [31:2]

MSI write address bits[31:2].

Bits [1:0]

Reads as 0b00.

Access to this field is RO.

Accessing the MPAMF_ERR_MSI_ADDR_L:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ERR_MSI_ADDR_L_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ERR_MSI_ADDR_L_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ERR_MSI_ADDR_L_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ERR_MSI_ADDR_L_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ERR_MSI_ADDR_L_s, MPAMF_ERR_MSI_ADDR_L_ns, MPAMF_ERR_MSI_ADDR_L_rt, and MPAMF_ERR_MSI_ADDR_L_rl must be separate registers:

- The Secure instance (MPAMF_ERR_MSI_ADDR_L_s) accesses the low part of the memory address for MSI write to signal an MPAM error used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ERR_MSI_ADDR_L_ns) accesses the low part of the memory address for MSI write to signal an MPAM error used for Non-secure PARTIDs.

- The Root instance (MPAMF_ERR_MSI_ADDR_L_rt) accesses the low part of the memory address for MSI write to signal an MPAM error used for Root PARTIDs.
- The Realm instance (MPAMF_ERR_MSI_ADDR_L_rl) accesses the low part of the memory address for MSI write to signal an MPAM error used for Realm PARTIDs.

MPAMF_ERR_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00E0	MPAMF_ERR_MSI_ADDR_L_s

Accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00E0	MPAMF_ERR_MSI_ADDR_L_ns

Accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00E0	MPAMF_ERR_MSI_ADDR_L_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ERR_MSI_ADDR_L can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00E0	MPAMF_ERR_MSI_ADDR_L_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.4 MPAMF_ERR_MSI_ATTR, MPAM Error MSI Write Attributes Register

The MPAMF_ERR_MSI_ATTR characteristics are:

Purpose

MPAMF_ERR_MSI_ATTR is a 32-bit read/write register that controls MPAM error MSI write attributes for MPAM errors in this MSC.

MPAMF_ERR_MSI_ATTR_s controls the attributes of Secure MPAM error MSI writes.

MPAMF_ERR_MSI_ATTR_ns controls the attributes of Non-secure MPAM error MSI writes.

MPAMF_ERR_MSI_ATTR_rt controls the attributes of Root MPAM error MSI writes.

MPAMF_ERR_MSI_ATTR_rl controls the attributes of Realm MPAM error MSI writes.

Configurations

The power domain of MPAMF_ERR_MSI_ATTR is IMPLEMENTATION DEFINED.


This register is present only when MPAMF_IDR.HAS_ERR_MSI == 1. Otherwise, direct accesses to MPAMF_ERR_MSI_ATTR are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ERR_MSI_ATTR is a 32-bit register.

Field descriptions



Bits [31:30]

Reserved, RES0.

MSI_SH, bits [29:28]

Sharability attribute of MSI writes.

- 0b00 Non-shareable.
- 0b01 Reserved, CONSTRAINED UNPREDICTABLE.
- 0b10 Outer Shareable.
- 0b11 Inner Shareable.

When MPAMF_ERR_MSI_ATTR.MSI_MEMATTR specifies a Device memory type, the contents of this field are IGNORED and Shareability is effectively Outer Shareable.

MSI_MEMATTR, bits [27:24]

Memory attributes of MSI writes.

Note: This encoding matches the VMSAv8-64 stage 2 MemAttr[3:0] field as described in the Arm ARM, except that the following encodings are Reserved (not UNPREDICTABLE) and behave as Device-nGnRnE: 0b0100, 0b1000, and 0b1100.

- 0b0000 Device-nGnRnE.
- 0b0001 Device-nGnRE.
- 0b0010 Device-nGRE.
- 0b0011 Device-GRE.
- 0b0100 Reserved. Behave as Device-nGnRnE, 0b0000.
- 0b0101 Normal Inner Non-cacheable, Outer Non-cacheable.
- 0b0110 Normal Inner Write-Through Cacheable, Outer Non-cacheable.

0b0111	Normal Inner Write-Back Cacheable, Outer Non-cacheable.
0b1000	Reserved. Behave as Device-nGnRnE, 0b0000.
0b1001	Normal Inner Non-Cachable, Outer Write-Through Cacheable.
0b1010	Normal Inner Write-Through Cacheable, Outer Write-Through Cacheable.
0b1011	Normal Inner Write-Back Cacheable, Outer Write-Through Cacheable.
0b1100	Reserved. Behave as Device-nGnRnE, 0b0000.
0b1101	Normal Inner Non-cacheable, Outer Write-Back Cacheable.
0b1110	Normal Inner Write-Through Cacheable, Outer Write-Back Cacheable.
0b1111	Normal Inner Write-Back Cacheable, Outer Write-Back Cacheable.

When this field specifies a Device memory type, the contents of MPAMF_ERR_MSI_ATTR.MSI_SH are IGNORED and Shareability is effectively Outer Shareable.

Device types may be implemented as any Device type with more than 'n' characters. For example, if this field is set to 0b0010, an implementation may treat the MSI write as the specified type, Device-nGRE, or as Device-nGnRE or as Device-nGnRnE.

Reserved encodings 0b0100, 0b1000, and 0b1100 must be implemented to behave the same as the 0b0000 encoding.

Bits [23:1]

Reserved, RES0.

MSIEN, bit [0]

Error interrupt MSI Enable.

0b0	MPAM error MSI writes are not generated to signal enabled MPAM error interrupts. When error MSI writes are disabled, hardwired error interrupts could be generated.
0b1	MPAM error MSI writes are generated to signal enabled MPAM error interrupts. When error MSI writes are enabled, hardwired error interrupts are not generated.

The value of this field affects whether hardwired error interrupts are generated.

The reset behavior of this field is:

- On a MSC reset, this field resets to 0.

Accessing the MPAMF_ERR_MSI_ATTR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ERR_MSI_ATTR_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ERR_MSI_ATTR_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ERR_MSI_ATTR_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ERR_MSI_ATTR_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ERR_MSI_ATTR_s, MPAMF_ERR_MSI_ATTR_ns, MPAMF_ERR_MSI_ATTR_rt, and MPAMF_ERR_MSI_ATTR_rl must be separate registers:

- The Secure instance (MPAMF_ERR_MSI_ATTR_s) accesses the memory access attributes for MSI write to signal an MPAM error used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ERR_MSI_ATTR_ns) accesses the memory access attributes for MSI write to signal an MPAM error used for Non-secure PARTIDs.

- The Root instance (MPAMF_ERR_MSI_ATTR_rt) accesses the memory access attributes for MSI write to signal an MPAM error used for Root PARTIDs.
- The Realm instance (MPAMF_ERR_MSI_ATTR_rl) accesses the memory access attributes for MSI write to signal an MPAM error used for Realm PARTIDs.

MPAMF_ERR_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00EC	MPAMF_ERR_MSI_ATTR_s

Accesses to this interface are RW.

MPAMF_ERR_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00EC	MPAMF_ERR_MSI_ATTR_ns

Accesses to this interface are RW.

MPAMF_ERR_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00EC	MPAMF_ERR_MSI_ATTR_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ERR_MSI_ATTR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00EC	MPAMF_ERR_MSI_ATTR_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.5 MPAMF_ERR_MSI_DATA, MPAM Error MSI Data Register

The MPAMF_ERR_MSI_DATA characteristics are:

Purpose

MPAMF_ERR_MSI_DATA is a 32-bit read/write register for the MPAM error MSI data.

MPAMF_ERR_MSI_DATA_s is the data for the MSI write for error interrupts related to Secure PARTIDs. MPAMF_ERR_MSI_DATA_ns is the data for the MSI write for error interrupts related to Non-secure PARTIDs. MPAMF_ERR_MSI_DATA_rt is the data for the MSI write for error interrupts related to Root PARTIDs. MPAMF_ERR_MSI_DATA_rl is the data for the MSI write for error interrupts related to Realm PARTIDs.

Configurations

The power domain of MPAMF_ERR_MSI_DATA is IMPLEMENTATION DEFINED.


This register is present only when MPAMF_IDR.HAS_ERR_MSI == 1. Otherwise, direct accesses to MPAMF_ERR_MSI_DATA are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ERR_MSI_DATA is a 32-bit register.

Field descriptions



MSI_DATA, bits [31:0]

MSI data to be written to ITS to signal an MSI.

Accessing the MPAMF_ERR_MSI_DATA:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ERR_MSI_DATA_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ERR_MSI_DATA_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ERR_MSI_DATA_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ERR_MSI_DATA_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ERR_MSI_DATA_s, MPAMF_ERR_MSI_DATA_ns, MPAMF_ERR_MSI_DATA_rt, and MPAMF_ERR_MSI_DATA_rl must be separate registers:

- The Secure instance (MPAMF_ERR_MSI_DATA_s) accesses the data for MSI write to signal an MPAM error used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ERR_MSI_DATA_ns) accesses the data for MSI write to signal an MPAM error used for Non-secure PARTIDs.
- The Root instance (MPAMF_ERR_MSI_DATA_rt) accesses the data for MSI write to signal an MPAM error used for Root PARTIDs.
- The Realm instance (MPAMF_ERR_MSI_DATA_rl) accesses the data for MSI write to signal an MPAM error used for Realm PARTIDs.

MPAMF_ERR_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00E8	MPAMF_ERR_MSI_DATA_s

Accesses to this interface are RW.

MPAMF_ERR_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00E8	MPAMF_ERR_MSI_DATA_ns

Accesses to this interface are RW.

MPAMF_ERR_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00E8	MPAMF_ERR_MSI_DATA_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ERR_MSI_DATA can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00E8	MPAMF_ERR_MSI_DATA_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.6 MPAMF_ERR_MSI_MPAM, MPAM Error MSI Write MPAM Information Register

The MPAMF_ERR_MSI_MPAM characteristics are:

Purpose

MPAMF_ERR_MSI_MPAM is a 32-bit read/write register that sets the MPAM information for error MSI write attributes for MPAM errors in this MSC.

MPAMF_ERR_MSI_MPAM_s controls MPAM information labeling of Secure MPAM error MSI writes. MPAMF_ERR_MSI_MPAM_ns controls MPAM information labeling of Non-secure MPAM error MSI writes. MPAMF_ERR_MSI_MPAM_rt controls MPAM information labeling of Root MPAM error MSI writes. MPAMF_ERR_MSI_MPAM_rl controls MPAM information labeling of Realm MPAM error MSI writes.

Configurations

The power domain of MPAMF_ERR_MSI_MPAM is IMPLEMENTATION DEFINED.


This register is present only when MPAMF_IDR.HAS_ERR_MSI == 1. Otherwise, direct accesses to MPAMF_ERR_MSI_MPAM are RES0.

The power and reset domain of each MSC component is specific to that component.

Attributes

MPAMF_ERR_MSI_MPAM is a 32-bit register.

Field descriptions



Bits [31:24]

Reserved, RES0.

PMG, bits [23:16]

Performance monitoring group property for PARTID MSC error interrupt write.

The reset behavior of this field is:

- On a MSC reset, this field resets to an architecturally UNKNOWN value.

PARTID, bits [15:0]

Partition ID for MSC error interrupt write.

The PARTID in this register is in the Secure PARTID space in the MPAMF_ERR_MSI_MPAM_s instance and in the Non-secure PARTID space in the MPAMF_ERR_MSI_MPAM_ns instance of this register.

The reset behavior of this field is:

- On a MSC reset, this field resets to an architecturally UNKNOWN value.

Accessing the MPAMF_ERR_MSI_MPAM:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ERR_MSI_MPAM_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ERR_MSI_MPAM_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ERR_MSI_MPAM_rt must only be accessible from the Root MPAM feature page.

- MPAMF_ERR_MSI_MPAM_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ERR_MSI_MPAM_s, MPAMF_ERR_MSI_MPAM_ns, MPAMF_ERR_MSI_MPAM_rt, and MPAMF_ERR_MSI_MPAM_rl must be separate registers:

- The Secure instance (MPAMF_ERR_MSI_MPAM_s) accesses the MPAM information for MSI write request to signal an MPAM error used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ERR_MSI_MPAM_ns) accesses the MPAM information for MSI write request to signal an MPAM error used for Non-secure PARTIDs.
- The Root instance (MPAMF_ERR_MSI_MPAM_rt) accesses the MPAM information for MSI write request to signal an MPAM error used for Root PARTIDs.
- The Realm instance (MPAMF_ERR_MSI_MPAM_rl) accesses the MPAM information for MSI write request to signal an MPAM error used for Realm PARTIDs.

MPAMF_ERR_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00DC	MPAMF_ERR_MSI_MPAM_s

Accesses to this interface are RW.

MPAMF_ERR_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00DC	MPAMF_ERR_MSI_MPAM_ns

Accesses to this interface are RW.

MPAMF_ERR_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00DC	MPAMF_ERR_MSI_MPAM_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ERR_MSI_MPAM can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00DC	MPAMF_ERR_MSI_MPAM_rl

When FEAT_RME is implemented, accesses to this interface are RW.

9.6.7 MPAMF_ESR, MPAM Error Status Register

The MPAMF_ESR characteristics are:

Purpose

Indicates MPAM error status for this MSC.

MPAMF_ESR_s reports Secure MPAM errors. MPAMF_ESR_ns reports Non-secure MPAM errors. MPAMF_ESR_rt reports Root MPAM errors. MPAMF_ESR_rl reports Realm MPAM errors.

Software should write this register after reading the status of an error to reset ERRCODE to 0x0000 and OVRWR to 0 so that future errors are not reported with OVRWR set.

Configurations

The power domain of MPAMF_ESR is IMPLEMENTATION DEFINED.

This register is present only when FEAT_MPAM is implemented. Otherwise, direct accesses to MPAMF_ESR are RES0.

MPAMF_ESR is 64-bit register when MPAM v0.1 or v1.1 is implemented and MPAMF_IDR.HAS_EXTD_ESR == 1.

Otherwise, MPAMF_ESR is a 32-bit register.

If an MSC cannot encounter any of the error conditions listed in [Errors in MSCs](#), both the MPAMF_ESR and MPAMF_ECR must be RAZ/WI.

The power and reset domain of each MSC component is specific to that component.


Attributes

MPAMF_ESR is a:

- 64-bit register when (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_IDR.HAS_EXTD_ESR == 1
- 32-bit register otherwise

Field descriptions

When (FEAT_MPAMv0p1 is implemented or FEAT_MPAMv1p1 is implemented) and MPAMF_IDR.HAS_EXTD_ESR == 1:



Bits [63:36]

Reserved, RES0.

RIS, bits [35:32]

When MPAMF_IDR.HAS_RIS == 1:

Resource Instance Selector. Where applicable to the ERRCODE, captures the RIS value for the error.

Otherwise:

Reserved, RES0.

OVRWR, bit [31]

Overwritten.

If 0 and ERRCODE == 0b0000, no errors have occurred.

If 0 and ERRCODE is non-zero, a single error has occurred and is recorded in this register.

If 1 and ERRCODE is non-zero, multiple errors have occurred and this register records the most recent error.

The state where this bit is 1 and ERRCODE is zero must not be produced by hardware and is only reached when software writes this combination into this register.

Bits [30:28]

Reserved, RES0.

ERRCODE, bits [27:24]

Error code.

- 0b0000 No error.
- 0b0001 PARTID_SEL_Range.
- 0b0010 Req_PARTID_Range.
- 0b0011 MSMONCFG_ID_RANGE.
- 0b0100 Req_PMG_Range.
- 0b0101 Monitor_Range.
- 0b0110 intPARTID_Range.
- 0b0111 Unexpected_INTERNAL.
- 0b1000 Undefined_RIS_PART_SEL.
- 0b1001 RIS_No_Control.
- 0b1010 Undefined_RIS_MON_SEL.
- 0b1011 RIS_No_Monitor.
- 0b1100 Reserved.
- 0b1101 Reserved.
- 0b1110 Reserved.
- 0b1111 Reserved.

PMG, bits [23:16]

Program monitoring group.

Set to the PMG on an error that captures PMG. Otherwise, set to 0x00 on an error that does not capture PMG.

PARTID_MON, bits [15:0]


PARTID or monitor.

Set to the PARTID on an error that captures PARTID.

Set to the monitor index on an error that captures MON.

On an error that captures neither PARTID nor MON, this field is set to 0.

Otherwise:



OVRWR, bit [31]

Overwritten.

If 0 and ERRCODE == 0b0000, no errors have occurred.

If 0 and ERRCODE is non-zero, a single error has occurred and is recorded in this register.

If 1 and ERRCODE is non-zero, multiple errors have occurred and this register records the most recent error.

The state where this bit is 1 and ERRCODE is 0 must not be produced by hardware and is only reached when software writes this combination into this register.

Bits [30:28]

Reserved, RES0.

ERRCODE, bits [27:24]

Error code.

0b0000	No error.
0b0001	PARTID_SEL_Range.
0b0010	Req_PARTID_Range.
0b0011	MSMONCFG_ID_RANGE.
0b0100	Req_PMG_Range.
0b0101	Monitor_Range.
0b0110	intPARTID_Range.
0b0111	Unexpected_INTERNAL.
0b1000	Reserved.
0b1001	Reserved.
0b1010	Reserved.
0b1011	Reserved.
0b1100	Reserved.
0b1101	Reserved.
0b1110	Reserved.
0b1111	Reserved.

PMG, bits [23:16]

Program monitoring group.

Set to the PMG on an error that captures PMG. Otherwise, set to 0x00 on an error that does not capture PMG.

PARTID_MON, bits [15:0]

PARTID or monitor.

Set to the PARTID on an error that captures PARTID.

Set to the monitor index on an error that captures MON.

On an error that captures neither PARTID nor MON, this field is set to 0x0000.

Accessing the MPAMF_ESR:

This register is within the MPAM feature page memory frames.

In a system that supports Secure, Non-secure, Root, and Realm memory maps, there must be MPAM feature pages in all four address maps:

- MPAMF_ESR_s must only be accessible from the Secure MPAM feature page.
- MPAMF_ESR_ns must only be accessible from the Non-secure MPAM feature page.
- MPAMF_ESR_rt must only be accessible from the Root MPAM feature page.
- MPAMF_ESR_rl must only be accessible from the Realm MPAM feature page.

MPAMF_ESR_s, MPAMF_ESR_ns, MPAMF_ESR_rt, and MPAMF_ESR_rl must be separate registers:

- The Secure instance (MPAMF_ESR_s) accesses the error status used for Secure PARTIDs.
- The Non-secure instance (MPAMF_ESR_ns) accesses the error status used for Non-secure PARTIDs.
- The Root instance (MPAMF_ESR_rt) accesses the error status used for Root PARTIDs.
- The Realm instance (MPAMF_ESR_rl) accesses the error status used for Realm PARTIDs.

MPAMF_ESR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_s	0x00F8	MPAMF_ESR_s

Accesses to this interface are RW.

MPAMF_ESR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_ns	0x00F8	MPAMF_ESR_ns

Accesses to this interface are RW.

MPAMF_ESR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rt	0x00F8	MPAMF_ESR_rt

When FEAT_RME is implemented, accesses to this interface are RW.

MPAMF_ESR can be accessed through its memory-mapped interface:

Component	Frame	Offset	Instance
MPAM	MPAMF_BASE_rl	0x00F8	MPAMF_ESR_rl

When FEAT_RME is implemented, accesses to this interface are RW.

Chapter 10

Errors in MSCs

This chapter contains the following sections:

- *Introduction.*
- *Error conditions in accessing memory-mapped registers.*
- *Overwritten error status.*
- *Behavior of configuration reads and writes with errors.*
- *Optionality of error detection and reporting.*

10.1 Introduction

When an MSC detects an error on an access to a memory-mapped register, information about the error must be captured in the MPAMF_ESR register and signaled to software through an interrupt. The errors covered by this mechanism might be caused by software errors.

Errors, whether detected or not, must not prevent the handling of the request by the MSC, but errors can cause the MSC to use different MPAM resource control settings than expected or cause monitors to misattribute monitored events. See *Optionality of error detection and reporting*.

Note

Implementation choices in an MSC might make certain errors impossible. For example, if the request interface only implements sufficient bits to exactly cover the range of 0 to PARTID_MAX and does not detect whether the unimplemented high-order bits of the PARTID are all zero, then the request PARTID cannot be detected as out-of-range, so ERRCODE == 2 cannot occur.

MPAM errors that an implementation detects are recorded in MPAMF_ESR_s or MPAMF_ESR_ns. The error condition descriptions in *Error conditions in accessing memory-mapped registers* describe whether the Security state of the PARTID or of the request address are used to determine which instance of MPAMF_ESR records the error status.

MSCs signal errors in accesses to memory-mapped registers using an error interrupt. See *MPAM Error Interrupt*. Errors recorded in MPAMF_ESR_s signal a Secure MPAM error interrupt if enabled by MPAMF_ECR_s.INTEN == 1. Errors recorded in MPAMF_ESR_ns signal a Non-secure MPAM error interrupt if enabled by MPAMF_ECR_ns.INTEN.

The MPAMF_ESR in an MSC captures the reason for an error, so that it can be accurately reported to software.

When *Resource instance selection* is implemented, hardware is permitted to make choices regarding CONSTRAINED UNPREDICTABLE behaviors and unimplemented RIS bits that could reduce or remove the need to detect or report any of the RIS-related errors. For more information on RIS, see *Resource instance selection*.

10.2 Error conditions in accessing memory-mapped registers

When an MSC detects an error condition, information about the error is captured in `MPAMF_ESR`. `MPAMF_ESR.ERRCODE` encodes the reason for the error as shown in Table 10-1. Other fields are captured in `MPAMF_ESR` as shown in the *Fields Captured* column of Table 10-1.

Table 10-1 Error conditions in accessing memory-mapped registers

MPAM Error Code (ERRCODE)	Error Name	Error Description	Fields Captured
0	No Error	No error captured in <code>MPAMF_ESR</code> .	None
1	<code>PARTID_SEL_Range</code>	<code>MPAMCFG_PART_SEL</code> stored with an out-of-range <code>PARTID</code> .	<code>PARTID</code> and <code>RIS</code> ^a
2	<code>Req_PARTID_Range</code>	A request has out-of-range <code>PARTID</code> .	<code>PARTID</code> , <code>PMG</code>
3	<code>MSMONCFG_ID_RANGE</code>	<code>MSMON</code> configuration request has out-of-range <code>PARTID</code> or <code>PMG</code> .	<code>PARTID</code> , <code>PMG</code> , <code>RIS</code> ^a
4	<code>Req_PMG_Range</code>	A request has out-of-range <code>PMG</code> .	<code>PARTID</code> and <code>PMG</code>
5	<code>Monitor_Range</code>	<code>MSMON_CFG_MON_SEL</code> has out-of-range monitor selector.	<code>MON_SEL</code> , <code>RIS</code> ^a
6	<code>intPARTID_Range</code>	The <code>intPARTID</code> in <code>MPAMCFG_INTPARTID</code> is out of the <code>intPARTID</code> range for the <code>PARTID</code> in <code>MPAMCFG_PART_SEL</code> .	<code>intPARTID</code>
7	<code>Unexpected_INTERNAL</code>	<code>MPAMCFG_PART_SEL.INTERNAL</code> is set when a <code>reqPARTID</code> is expected.	<code>PARTID</code>
8	<code>Undefined_RIS_PART_SEL</code>	Unimplemented <code>RIS</code> in <code>MPAMCFG_PART_SEL.RIS</code> .	<code>PART_SEL</code> , <code>RIS</code>
9	<code>RIS_No_Control</code>	Resource instance selected by <code>MPAMCFG_PART_SEL.RIS</code> does not have the accessed partitioning control.	<code>PART_SEL</code> , <code>RIS</code>
10	<code>Undefined_RIS_MON_SEL</code>	Unimplemented <code>RIS</code> in <code>MSMON_CFG_MON_SEL</code> .	<code>MON_SEL</code> , <code>RIS</code>
11	<code>RIS_No_Monitor</code>	Resource instance selected by <code>MSMON_CFG_MON_SEL.RIS</code> does not have the accessed monitor type.	<code>MON_SEL</code> , <code>RIS</code>
12:18	Reserved	Reserved for future use.	--

a. This field is only available when `MPAMF_IDR.EXT` and `MPAMF_IDR.HAS_RIS` are 1.

10.2.1 No error (errorcode == 0)

No error is captured in `MPAMF_ESR`.

10.2.2 `PARTID_SEL` out-of-range error (errorcode == 1)

The value of the `MPAMCFG_PART_SEL.PARTID_SEL` field is out-of-range for the `PARTID` space selected by the `NS` bit on a store to an `MPAMCFG` memory-mapped register.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is also controlled by the NS bit.

10.2.3 Request PARTID out-of-range error (errorcode == 2)

The value of PARTID in a request is out-of-range for the MSC's MPAM implementation of PARTID space selected by the MPAM_NS bit.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is also controlled by the MPAM_NS bit.

The MPAM behavior of an MSC for a request that causes this error is CONstrained UNPREDICTABLE:

- The request may be processed as if the PARTID is any valid PARTID in the same MPAM Security state (MPAM_NS) as the request's PARTID.
- Arm recommends that the default PARTID for the MPAM_NS Security state is used.

10.2.4 MSMON configuration ID out-of-range error (errorcode == 3)

A write to configure a monitor contains an out-of-range value for either the PARTID or PMG for the PARTID space selected by the Secure address space bit, NS.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is also controlled by the NS bit.

10.2.5 Request PMG out-of-range error (errorcode == 4)

The value of PMG in a request is out of range for the MSC's MPAM implementation of the PMG space selected by the MPAM security space bit, MPAM_NS.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is also controlled by the MPAM_NS bit.

The MPAM behavior of an MSC for a request that causes this error is CONstrained UNPREDICTABLE:

- The request may be processed as if the PARTID and PMG are any valid PARTID and PMG in the same MPAM Security state as the request.
 - Arm recommends that the request be processed as if the PMG is the default.
- The default PARTID and PMG may be used for the request's MPAM_NS Security state. The request may be IGNORED for performance monitoring, as if the PMG value does not match the monitor's PMG filter even if the PARTID matches.

10.2.6 Monitor out-of-range error (errorcode == 5)

The value of the monitor selector register, `MSMON_CFG_MON_SEL.MON_SEL`, is out of range on a store to an `MSMON_*` memory-mapped register selected by the Secure address space bit, NS.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is also controlled by the NS bit.

10.2.7 intPARTID out-of-range error (errorcode == 6)

This error can only occur if PARTID narrowing is implemented. `MPAMF_IDR.HAS_PARTID_NRW == 1` indicates that an implementation has PARTID narrowing.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is controlled by the Secure address space bit, NS.

These conditions cause this error:

- `MPAMCFG_INTPARTID.INTPARTID` is out-of-range for the `intPARTID` space selected by the Secure address space bit, NS, on a store to a memory-mapped register to configure the association of `reqPARTID` to `intPARTID`.
- `MPAMCFG_INTPARTIDINTERNAL == 0` on any write to configure `MPAMCFG_INTPARTID`.
- `MPAMCFG_PART_SEL.INTERNAL` is not set when an `intPARTID` is expected. These expected cases include a read or write to any `MPAMCFG_*` register, other than `MPAMCFG_INTPARTID`.

10.2.8 Unexpected INTERNAL error (errorcode == 7)

This error can only occur if PARTID narrowing is implemented. `MPAMF_IDR.HAS_PARTID_NRW == 1` indicates that an implementation has PARTID narrowing.

If PARTID narrowing is implemented in the MSC, this error is detected if the `MPAMCFG_PART_SEL.INTERNAL` bit is set when a `reqPARTID` is expected. When PARTID narrowing is implemented, the only cases in which a `reqPARTID` is expected in `MPAMCFG_PART_SEL` are a read or write access to `MPAMCFG_INTPARTID`.

The selection of the Secure or Non-secure version of `MPAMF_ESR` for capturing the error information is controlled by the Secure address space bit, NS.

Reads that cause this error return an UNKNOWN value.

10.2.9 Undefined RIS in MPAMCFG_PART_SEL.RIS (errorcode == 8)

This error occurs when an access to an `MPAMCFG_*` register occurs when `MPAMCFG_PART_SEL.RIS` does not correspond to a RIS value allocated to an MPAM resource of the MSC. The MPAM behavior of an MSC for a request that causes this error is a CONSTRAINED UNPREDICTABLE choice between:

- RAZ/WI.
- RAZ/WI and record an MPAM error in the `MPAMF_ESR` associated with that MSC, using the error code `ERRCODE == 8` and capturing `MPAMCFG_PART_SEL.{RIS, PARTID_SEL}`.

10.2.10 RIS in MPAMCFG_PART_SEL.RIS does not have partitioning control (errorcode == 9)

This error occurs when an access to an `MPAMCFG_*` register occurs when `MPAMCFG_PART_SEL.RIS` selects a resource that exists but does not have the partitioning control accessed. The MPAM behavior of an MSC for a request that causes this error is a CONSTRAINED UNPREDICTABLE choice between:

- RAZ/WI.
- RAZ/WI and record an MPAM error in the `MPAMF_ESR` associated with that MSC, using the error code `ERRCODE == 9` and capturing `MPAMCFG_PART_SEL.{RIS, PARTID_SEL}`.

10.2.11 Undefined RIS in MSMON_CFG_MON_SEL.RIS (errorcode == 10)

This error occurs when an access to an `MSMON_CFG_*` register occurs when `MSMON_CFG_MON_SEL.RIS` does not correspond to an MPAM resource of the MSC. The MPAM behavior of an MSC for a request that causes this error is a CONSTRAINED UNPREDICTABLE choice between:

- RAZ/WI.
- RAZ/WI and record an MPAM error in the `MPAMF_ESR` associated with that MSC, using the error code `ERRCODE == 10` and capturing `MSMON_CFG_MON_SEL.{RIS, MON_SEL}`.

10.2.12 RIS selected by MSMON_CFG_MON_SEL.RIS does not have monitor type (errorcode == 11)

Access to an MSMON_<type> or MSMON_<type>_CAPTURE register when MSMON_CFG_MON_SEL.RIS does not correspond to an MPAM resource of the MSC or that does not have the type of monitor accessed by the MSMON_<type> or MSMON_<type>_CAPTURE register. The MPAM behavior of an MSC for a request that causes this error is a CONSTRAINED UNPREDICTABLE choice between:

- Read as 0xFFFFFFFF, NRDY == 1 with value of 0x7FFFFFFE, and WI. This value is highly unlikely as a normal return value in any monitor.
- RAZ/WI.
- RAZ/WI and record an MPAM error in the MPAMF_ESR associated with that MSC, using the error code ERRCODE == 11 and capturing MSMON_CFG_MON_SEL.{RIS, MON_SEL}.

Access to an MSMON_<type>_* register when MSMON_CFG_MON_SEL.RIS does not correspond to an MPAM resource that has the type of monitor accessed by the MSMON_<type>_* register is CONSTRAINED UNPREDICTABLE, one of:

- RAZ/WI.
- RAZ/WI and record an MPAM error in the MPAMF_ESR associated with that MSC, using the error code ERRCODE == 11 and capturing MSMON_CFG_MON_SEL.{RIS, MON_SEL}.

10.2.13 Reserved (errcodes 12 – 15)

These error codes are reserved for future use.

10.3 Overwritten error status

When MPAMF_ESR is written due to an error, and the ERRCODE field was not previously 0, the OVRWR bit is set. Error status is always written to MPAMF_ESR, whether or not it contains a previously recorded error syndrome.

Table 10-2 Overwritten error status

OVRWR	ERRCODE	Description
0	0b0000	No errors have been recorded in MPAMF_ESR.
0	Non-zero	Not overwritten. A single error has been written to MPAMF_ESR since it was last cleared.
1	0b0000	This state is not produced by hardware, only by a software write.
1	Non-zero	Overwritten. Two or more errors have been written to MPAMF_ESR with only the syndrome information from the latest error recorded into the fields.

The interrupt service routine should clear both the ERRCODE and OVRWR fields of MPAMF_ESR after its contents have been read. This allows the OVRWR bit to accurately indicate when one or more errors have been overwritten before servicing future MPAM error interrupts.

10.4 Behavior of configuration reads and writes with errors

10.4.1 Writing an out-of-range PARTID to MPAMCFG_PART_SEL.PARTID_SEL

If a write to `MPAMCFG_PART_SEL` has a `PARTID_SEL` value that is out-of-range, it is IMPLEMENTATION DEFINED whether:

- The contents written to `MPAMCFG_PART_SEL.PARTID_SEL` are not checked at the time of the write and store the new value into `MPAMCFG_PART_SEL.PARTID_SEL`. The written out-of-range value could later cause a `PARTID_SEL` out-of-range error (`ERRCODE = 1`) when used to index an access to another configuration register by `PARTID_SEL`. See *Required error condition detection* for more information about the optionality of error detection.
- The contents being written to `MPAMCFG_PART_SEL.PARTID_SEL` are checked before updating the `MPAMCFG_PART_SEL` register. If the error is detected, the `MPAMCFG_PART_SEL` register is not updated and the `PARTID_SEL` out-of-range error (`ERRCODE = 1`) is raised. To implement this behavior, the implementation must detect the error.

10.4.2 Reading another MPAMCFG_* register when MPAMCFG_PART_SEL.PARTID_SEL contains an out-of-range PARTID

A read of any `MPAMCFG_*` register other than `MPAMCFG_PART_SEL` when `MPAMCFG_PART_SEL.PARTID_SEL` contains an out-of-range `PARTID` raises a `PARTID_SEL` out-of-range error (`ERRCODE = 1`) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

It is IMPLEMENTATION DEFINED whether the value returned by a read of another `MPAMCFG_*` register when `MPAMCFG_PART_SEL.PARTID_SEL` contains an out-of-range `PARTID` that is detected:

- Is an UNKNOWN value.
- Is a constant value of zero in all fields.

The value returned by a read of another `MPAMCFG_*` register when `MPAMCFG_PART_SEL.PARTID_SEL` contains an out-of-range `PARTID` that is not detected is an UNKNOWN value.

———— Note —————

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect out-of-range `PARTID_SEL` values and to not update the `MPAMCFG_PART_SEL` register, it is not possible to have an out-of-range `PARTID_SEL` value in that register and the precondition for this section cannot occur. See *Writing an out-of-range PARTID to MPAMCFG_PART_SEL.PARTID_SEL*.

10.4.3 Writing another MPAMCFG_* register when MPAMCFG_PART_SEL.PARTID_SEL contains an out-of-range PARTID

A write of any `MPAMCFG_*` register other than `MPAMCFG_PART_SEL` when `MPAMCFG_PART_SEL.PARTID_SEL` contains an out-of-range `PARTID` raises a `PARTID_SEL` out-of-range error (`ERRCODE = 1`) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

If a write to an `MPAMCFG_*` register other than `MPAMCFG_PART_SEL` has a `PARTID_SEL` out-of-range error (`ERRCODE = 1`), whether that error is detected or not detected, it is IMPLEMENTATION DEFINED whether:

- The write updates the configuration register indexed by an UNKNOWN in-range `PARTID`.
- The write is ignored (WI).

Note

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect out-of-range PARTID_SEL values and to not update the MPAMCFG_PART_SEL register, it is not possible to have an out-of-range PARTID_SEL value in that register and the precondition for this section cannot occur. See *Writing an out-of-range PARTID to MPAMCFG_PART_SEL.PARTID_SEL*.

10.4.4 Writing an undefined RIS to MPAMCFG_PART_SEL.RIS

If RIS is implemented and a configuration write to MPAMCFG_PART_SEL.RIS has an Undefined RIS error (ERRCODE = 8), it is IMPLEMENTATION DEFINED whether:

- The contents written to MPAMCFG_PART_SEL.RIS are not checked at the time of the write and store the new value in MPAMCFG_PART_SEL.RIS. This undefined RIS value could cause an Undefined RIS error (ERRCODE = 8) when later used to select a resource on an access to a configuration register by PARTID_SEL and RIS.
- The contents being written to MPAMCFG_PART_SEL.RIS are checked before updating the MPAMCFG_PART_SEL register. If the error is detected, the MPAMCFG_PART_SEL register is not updated and the Undefined RIS error (ERRCODE = 8) is raised. To implement this behavior, the implementation must detect the error.

10.4.5 Reading other MSC MPAM registers when MPAMCFG_PART_SEL.RIS contains an undefined RIS value

A read of an MPAMF*IDR register or an MPAMCFG_* register other than MPAMCFG_PART_SEL when MPAMCFG_PART_SEL.RIS contains an undefined RIS value raises an Undefined RIS error (ERRCODE = 8) if the implementation detects that error. See *Required error condition detection* for more information about the optionality of error detection. If the error is not detected, the value returned is UNKNOWN.

The value read from an MPAMF*IDR or an MPAMCFG_* register other than MPAMCFG_PART_SEL when MPAMCFG_PART_SEL.RIS contains a RIS value that does not correspond to an implemented resource instance returns an UNKNOWN value.

Note

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect undefined RIS values and to not update the MPAMCFG_PART_SEL register, it is not possible to have an undefined RIS value in that register and the precondition for this section cannot occur. See *Writing an undefined RIS to MPAMCFG_PART_SEL.RIS*.

10.4.6 Writing other MSC MPAM registers when MPAMCFG_PART_SEL.RIS contains an undefined RIS value

A write of an MPAMCFG_* register other than MPAMCFG_PART_SEL when MPAMCFG_PART_SEL.RIS contains an undefined RIS value raises an Undefined RIS error (ERRCODE = 8) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

If a configuration write to an MPAMCFG_* register other than MPAMCFG_PART_SEL has a RIS value that does not correspond to an implemented resource instance, whether the undefined RIS error is detected or not detected, it is IMPLEMENTATION DEFINED whether:

- The write might update the configuration register for any implemented resource instance.
- The write is ignored (WI).

Note

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect undefined RIS values and to not update the MPAMCFG_PART_SEL register, it is not possible to have an undefined RIS value in that register and the precondition for this section cannot occur. See *Writing an undefined RIS to MPAMCFG_PART_SEL.RIS*.

10.4.7 Reads of MSC MPAM registers with other errors

If there is no PARTID_SEL out-of-range error (ERRCODE = 1) and no Undefined RIS error (ERRCODE = 8), a configuration read to an MPAM*IDR or an MPAMCFG_* register that has any other errors detected returns an UNKNOWN value.

10.4.8 Writes to MSC MPAM registers with other errors

If there is no PARTID_SEL out-of-range error (ERRCODE = 1) and no Undefined RIS error (ERRCODE = 8), a configuration write to an MPAMCFG_* register that has any other errors detected leaves the control settings for the partition selected by MPAMCFG_PART_SEL.PARTID_SEL and MPAMCFG_PART_SEL.RIS in an UNKNOWN state.

10.4.9 Writes to MSMON_CFG_MON_SEL.MON_SEL

Writes to MSMON_CFG_MON_SEL that have the MON_SEL field out-of-range for the monitors of the MSC cannot generally be detected when the MON_SEL register is written because different types of monitors could have different numbers of supported monitor instances. If RIS is also implemented, then the resource instance selector being written into the RIS field could change which monitor types are available and how many monitor instances of each type are implemented because different resource instances could have different numbers of monitor instances from the same resource type.

There are limited cases where MSMON_CFG_MON_SEL.MON_SEL could be checked when written:

- RIS is not implemented and only a single monitor type is supported.
- RIS is not supported and all supported monitor types have exactly the same number of monitor instances.
- RIS is supported and all monitor types of all resource instances support exactly the same number of monitor instances.
- RIS is supported, different resource instances support a different number of monitor instances, and all monitor types of each resource instance support exactly the same number of monitor instances. In this case the RIS value must be used to determine the maximum number of monitor instances to check the MON_SEL value.

Checking for out-of-range MON_SEL when MSMON_CFG_MON_SEL is written is an implementation option because some of the detectable cases could be common.

If a configuration write to MSMON_CFG_MON_SEL has a MON_SEL value that is out-of-range, it is IMPLEMENTATION DEFINED whether:

- The contents written to MSMON_CFG_MON_SEL.MON_SEL are not checked at the time of the write and store the new value into the register. The written out-of-range value could later cause a MON_SEL out-of-range error (ERRCODE = 5) when used to index an access to a MSMON_CFG_* configuration register or MSMON_* monitor or capture register by MON_SEL.
- The contents being written to MSMON_CFG_MON_SEL.MON_SEL are checked before updating the MSMON_CFG_MON_SEL register. If the error is detected, the MSMON_CFG_MON_SEL register is not updated and the MON_SEL out-of-range error (ERRCODE = 5) is raised. See *Required error condition detection* for more information about the optionality of error detection.

10.4.10 Reading another MSMON_* register when MSMON_CFG_MON_SEL.MON_SEL out of range

A read of any MSMON_* register other than MSMON_CFG_MON_SEL when MSMON_CFG_MON_SEL.MON_SEL contains an out-of-range monitor instance selector raises a Monitor Range error (ERRCODE = 5) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

The value read from any MSMON_* register other than MSMON_CFG_MON_SEL when MSMON_CFG_MON_SEL.MON_SEL contains an out-of-range monitor instance selector returns an UNKNOWN value whether the Monitor Range error is detected or not detected.

———— **Note** —————

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect out-of-range MON_SEL values and to not update the `MSMON_CFG_MON_SEL` register, it might not be possible to have an out-of-range MON_SEL value in that register and the precondition for this section cannot occur. Section *Writes to `MSMON_CFG_MON_SEL.MON_SEL`* lists the conditions necessary to permit the choice of this option.

10.4.11 Writes to `MSMON_*` registers with `MSMON_CFG_MON_SEL.MON_SEL` out of range

A write of any `MSMON_*` register other than `MSMON_CFG_MON_SEL` when `MSMON_CFG_MON_SEL.MON_SEL` contains an out-of-range monitor instance selector, raises a Monitor Range error (ERRCODE = 5) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

If a write is to an `MSMON_*` register other than `MSMON_CFG_MON_SEL` when `MSMON_CFG_MON_SEL.MON_SEL` is out-of-range, whether the error is detected or not detected, it is IMPLEMENTATION DEFINED whether:

- The write could update an `MSMON_*` register indexed by any in-range monitor instance selector.
- The write is ignored (WI).

———— **Note** —————

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect out-of-range MON_SEL values and to not update the `MSMON_CFG_MON_SEL` register, it might not be possible to have an out-of-range MON_SEL value in that register and the precondition for this section cannot occur. *Writes to `MSMON_CFG_MON_SEL.MON_SEL`* lists the conditions necessary to permit the choice of this option.

10.4.12 Writing an undefined RIS to `MSMON_CFG_MON_SEL.RIS`

If RIS is implemented and a configuration write to `MSMON_CFG_MON_SEL.RIS` has a value that does not correspond to an implemented resource instance, it is IMPLEMENTATION DEFINED whether:

- The value written to `MSMON_CFG_MON_SEL.RIS` is not checked at the time of the write and the new values are stored in that register. This undefined RIS value could cause an Undefined_RIS_MON_SEL error (ERRCODE = 10) when later used to select a resource on an access to an `MSMON_*` register by `MON_SEL` and `RIS`.
- The contents being written to `MSMON_CFG_MON_SEL.RIS` are checked before updating the `MSMON_CFG_MON_SEL` register. If the error is detected, the register is not updated and the Undefined_RIS_MON_SEL error (ERRCODE = 10) is raised.

10.4.13 Reading another `MSMON_*` register when `MSMON_CFG_MON_SEL.RIS` contains an undefined RIS value

A read of an `MSMON_*` register other than `MSMON_CFG_MON_SEL` when `MSMON_CFG_MON_SEL.RIS` contains a RIS value that does not correspond to an implemented resource instance raises an Undefined_RIS_MON_SEL error (ERRCODE = 10) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

The value read from an `MSMON_*` register other than `MSMON_CFG_MON_SEL` when `MSMON_CFG_MON_SEL.RIS` contains a RIS value that does not correspond to an implemented resource instance returns an UNKNOWN value whether the error is detected or not detected.

———— **Note** —————

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect undefined RIS values and to not update the `MSMON_CFG_MON_SEL` register, it is not possible to have an undefined RIS value in that register and the precondition for this section cannot occur. See *Writing an undefined RIS to `MSMON_CFG_MON_SEL.RIS`*.

10.4.14 Writing another MSMON_* register when MSMON_CFG_MON_SEL.RIS contains an undefined RIS value

A write of an MSMON_* register other than MSMON_CFG_MON_SEL when MSMON_CFG_MON_SEL.RIS contains a RIS value that does not correspond to an implemented resource instance raises an Undefined_RIS_MON_SEL error (ERRCODE = 10) if that error is detected. See *Required error condition detection* for more information about the optionality of error detection.

If a write to an MSMON_* register other than MSMON_CFG_MON_SEL has a RIS value that does not correspond to an implemented resource, whether the undefined RIS error is detected or not detected, it is IMPLEMENTATION DEFINED whether:

- The write might update the MSMON_* register indexed by any implemented resource instance.
- The write is ignored (WI).

———— **Note** ————

In an implementation that chooses the IMPLEMENTATION DEFINED option to detect undefined RIS values and to not update the MSMON_CFG_MON_SEL register, it is not possible to have an undefined RIS value in that register and the precondition for this section cannot occur. See *Writing an undefined RIS to MSMON_CFG_MON_SEL.RIS*.

—————

10.5 Optionality of error detection and reporting

Error detection and reporting are required for an error condition when all of the following are true:

- The MSC supports at least one MPAM feature that can raise the error condition.
- The MSC is designed so that the particular error condition can occur.
- The MSC is required to detect the error condition, see *Required error condition detection*.

If there are no error conditions that meet these criteria, then in MPAM v0.1 and from MPAM v1.1, `MPAMF_IDR.HAS_ESR` is permitted to be 0. If `MPAMF_IDR.HAS_ESR` is 1, then `MPAMF_ESR` and `MPAMF_ECR` must be implemented.

In MPAM v1.0, if no error conditions are detected, `MPAMF_ESR` and `MPAMF_ECR` must be RAZ/WI.

10.5.1 Required error condition detection

This section describes the conditions under which each of the MPAM MSC error conditions must be detected. In cases where detection is not required, an implementation might choose not to implement detection and reporting logic for that error condition.

10.5.1.1 Selector out-of-range errors

The following requirements apply to each of the types of selectors used in MPAM in MSCs, including:

- PARTID.
- PMG.
- Monitor selectors.
- In MPAM v0.1 and from MPAM v1.1, RIS values.

The selector interface is permitted to be narrower than the full width specified in the architecture. Even if the MSC interface is of one size, the internal implementation might be smaller than that size. Bits beyond the implemented width of any selector are permitted to be silently truncated without any requirement to detect or report should those bits be non-zero.

An MSC implementation that supports a range that is not 0 to $2^n - 1$ in a field of n bits for any selector is required to detect and report values that lie within the field size but are not valid in the implementation. Such detection can be applied after performing the silent truncation to the bit-width supported.

10.5.1.2 PARTID narrowing errors

If PARTID narrowing is supported, the Unexpected Internal error condition must be detected and reported.

Appendix A

Generic Resource Controls

This chapter contains the following sections:

- *Introduction.*
- *Portion resource controls.*
- *Maximum-usage resource controls.*
- *Proportional resource allocation facilities.*
- *Combining resource controls.*

A.1 Introduction

This appendix is *Informative*.

Several of the resource controls defined in this specification fit one of the generic models for resource controls in this appendix.

A.2 Portion resource controls

Some resources may be divided into fixed quanta, termed *portions*, that can be allocated for the exclusive use of a partition or shared between two or more partitions. Figure A-1 shows how partitions can have private and shared Portion Bit Map (PBM) allocations.

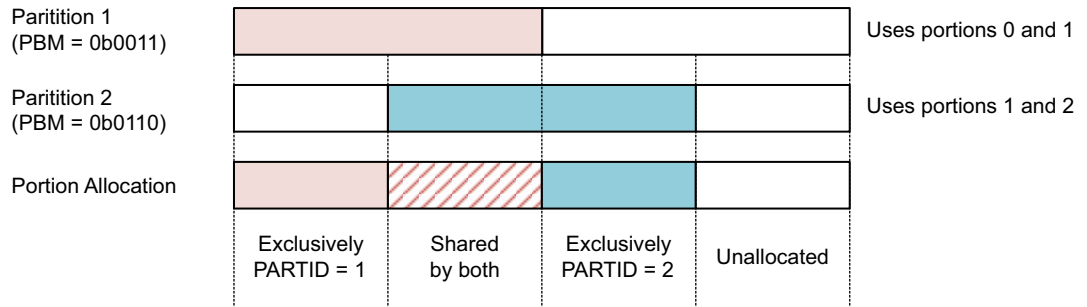


Figure A-1 Generic portion shared and exclusive allocations.

In portion resource controls, the control setting is a bitmap in which each bit corresponds to a particular portion of the resource, as shown in Figure A-2. Each bit grants the PARTID using this control setting to allocate the portion corresponding to that bit.




Figure A-2 Generic portion bit map.

PBMs can be wide. Generic PBMs can be up to 2^{15} bits in width.

A PBM is a vector of single-bit elements. Element 0 is bit 0 at the address $(MPAMF_BASE + PBM_offset)$ where PBM_offset is the offset of the particular PBM register. Both the bitmap and the register to access the bitmap extend in length at increasing 32-bit word addresses for the width in bits of the PBM (PBM_WD). If the 32-bit word containing the highest byte of the bitmap $(MPAMF_BASE + PBM_offset + (PBM_WD \gg 3))$ has unused bits, those bits are RES0.

To access the PBM for portion n , access the 32-bit word of the PBM register at the address $MPAMF_BASE + PBM_offset + ((n \gg 3) \& \sim 3)$. Then access bit $(n \& 31)$.

A.3 Maximum-usage resource controls

Many resources can be controlled by a maximum-usage resource control. With this control, resources may be allocated to a partition as long as the partition's maximum usage is not exceeded. If the maximum usage is reached, further allocation must be prevented, or deferred, or lowered in priority, or caused to reclaim a previous allocation, or caused to replace a previous allocation.

Maximum-usage control settings are a maximum fraction of the resource that the PARTID may use. The parameter is represented as a 16-bit fixed-point fraction of the capacity of the resource with a discoverable number of fractional bits. For example, if a resource has an 8-bit fractional width, bits [15:8] of the setting are used to control the resource allocation. To ensure that the range includes 100% of the resource, the control value is increased by 1 in the least significant implemented bit before being used to limit the usage to the maximum. See [About the fixed-point fractional format](#) for the fixed-point fractional format.

A.4 Proportional resource allocation facilities

MPAM proportional stride partitioning is related to two software resource-management interfaces:

- The Linux cgroup weights interface assigns integer weights to indicate the relative proportion of the resource given to a process.
- The VMware shares interface similarly assigns an integer share to indicate the relative proportion of the resource that a virtual machine is given.

Weight and share values are positive integers. For example, Linux group weights are in the range of 1 to 10000, with a default value of 100.

The value of weight or share is used to compute the fraction of the resource, f , for partition, p , as:

$$f(p) = \frac{\text{Weight}_p}{\sum_{\text{all } w} \text{Weight}_w}$$

A partition's stride is the scaled reciprocal of its weight:

$$\text{Stride of } p = \frac{S}{f(p)}$$

The scaling factor, S , should be chosen as equal to the largest $f(p)$ so as to normalize stride values and give the smallest stride in the system = 1. All strides should be scaled by the same S .

Stride-based proportional allocation is well-suited to temporal or rate-of-occurrence resources, such as bandwidth.

The standard interface for proportional allocation is a positive unsigned integer, STRIDEM1, with an IMPLEMENTATION DEFINED field width of w . STRIDEM1 has the range $[0 \dots 2^w - 1]$ so stride has the range $[1 \dots 2^w]$. If a stride after normalization is greater than 2^w , it should be programmed into the control as $2^w - 1$, the largest representable STRIDEM1.

Properties of proportional allocation include:

- Proportion of resource shrinks and grows as partitions come and go.
- Subdividable: If VM A has $\frac{1}{2}$ fraction of the whole resource and its child application, y, has $\frac{2}{3}$ fraction of the VM's resource, then y is given $\frac{1}{2} * \frac{2}{3} = \frac{1}{3}$ fraction of the whole resource.
- Proportional allocation only needs to consider the current contenders for a temporal resource, such as memory bandwidth.
- A proportional allocation scheme is called *work-conserving* if it does not idle the resource when only low-proportion requests are available, but instead uses as much of the resource as it has requests to use. A proportional allocation scheme might allocate the resource to those lower-proportion requests, in proportion to their relative weights.

A.4.1 Model of stride-based memory bandwidth scheduling

This model is intended to explain the operation of stride-based memory bandwidth scheduling without dictating an implementation. Arm believes that a variety of implementations are possible.

In this model, each partition has an $offset[p]$ that tracks the time since the partition, p , consumed bandwidth but is bounded to be less than $offset_limit$. When a request, r , arrives it is given a *deadline*, of the *current_time* plus $stride(p)$ minus $offset(p)$. The $offset(p)$ is set to *current_time* – *deadline*, and the $offset(p)$ is incremented in event-time units until it reaches the $offset_limit$.

In the model, requests are serviced as quickly as possible in deadline order. Newly arriving requests with small strides (highest access to bandwidth) may go ahead of earlier requests with large strides.

If there are requests to process, this model does not prevent servicing a request with a distant future deadline if there are no requests available with earlier deadlines. As such, this model scheme is work-conserving.

A.5 Combining resource controls

Maximum-usage resource controls, portion resource controls, and other resource controls may coexist on the same resource. Combined resource controls should produce a combined effect. For example, combining portion control and maximum-usage control for the same resource should allocate the resource while satisfying both controls.

All resource controls should have at least one setting that does not limit access to the resource. When an implementation contains multiple controls for the same resource, the limits imposed on a partition's usage by each control are all applied. By selecting which controls limit a partition's usage and which do not, software can exercise a variety of regulation styles within a single system.

Appendix B

MSC Firmware Data

This chapter contains the following sections:

- *Introduction.*
- *Partitioning-control parameters.*
- *Performance-monitoring parameters.*
- *Discovery of resource to RIS mapping.*
- *Discovery of wired interrupts.*

B.1 Introduction

In a system containing MPAM, discovery of the memory-system topology and certain implementation parameters of MPAM controls and monitors must be provided to MPAM-aware software via firmware data. The software-to-firmware interface to the MPAM firmware data is beyond the scope of this description. Examples of firmware data interfaces include:

- ACPI.
- Device Tree.

Firmware data for static devices can be pre-configured for an implementation and stored as part of the firmware, or it can be dynamically discovered through probing and other tests, or some combination of these two approaches.

B.2 Partitioning-control parameters

Table B-1 Partitioning-control parameters.

Control	Parameter	Data Format	Description
MPAM	MPAMF_BASE_NS	Address	Every MPAM-capable device has the MPAMF_IDR MMR at offset 0 from the MPAMF_BASE_NS in the Non-secure address space. Other MPAM memory-mapped registers are at known offsets from this address. See Chapter 9 Memory-mapped Registers .
MPAM	MPAMF_BASE_S	Address	Every MPAM-capable device has the MPAMF_IDR MMR at offset 0 from the MPAMF_BASE_S in the Secure address space. Other MPAM memory-mapped registers are at known offsets from this address. See Chapter 9 Memory-mapped Registers .

B.3 Performance-monitoring parameters

Table B-2 Performance-monitoring parameters

Monitor	Parameter	Data Format	Description
CSU	MAX_NRDY_USEC	Uint32	Maximum number of microseconds that the NRDY signal can remain 1 in the absence of additional reconfiguration of the monitor or writes to the MSMON_CSU register. This firmware value is the maximum time when NRDY can be 1, so that software can know this value. MSMON_CSU.VALUE is accurate and MSMON_CSU.NRDY is zero before MAX_NRDY_USEC microseconds have elapsed since the monitor was configured, reconfigured, or written.

B.4 Discovery of resource to RIS mapping

Software needs to know which RIS value to use to control a resource instance of the MSC.

This mapping is not available from MSC IDRs. It might be given as a firmware data table or other means beyond the hardware ID registers.

B.5 Discovery of wired interrupts

There are two interrupt sources in an MPAM MSC and they are replicated in the Secure and Non-secure MPAM behaviors. It is not possible to discover the connection of the four interrupts to GIC inputs from the MSC MPAM ID registers. This information must come from the firmware information.

Firmware must provide information on the connection and grouping of MPAM wired interrupts.

Glossary

This glossary describes some of the terms that are used in this document. Some of these terms are unique to MPAM and are introduced in this document while others are standard terms that can be found in the Glossary of the *Arm® Architecture Reference Manual for A-profile architecture*.

Abort	An exception caused by an illegal memory access. Aborts can be caused by the external memory system or the MMU.
Aligned	A data item stored at an address that is exactly divisible by the highest power of 2 that divides exactly into its size in bytes. Aligned halfwords, words and doublewords therefore have addresses that are divisible by 2, 4, and 8, respectively.
ALTSP	Alternative PARTID space.
AMBA	Advanced Microcontroller Bus Architecture. The AMBA family of protocol specifications is the Arm open standard for on-chip buses. AMBA provides solutions for the interconnection and management of the functional blocks that make up a <i>System-on-Chip</i> (SoC). Applications include the development of embedded systems with one or more processors or signal processors and multiple peripherals.
Banked register	A register that has multiple instances, with the instance that is in use depending on the PE mode, Security state, or other PE state.

Burst A group of transfers that form a single transaction. With AMBA protocols, only the first transfer of the burst includes address information, and the transfer type determines the addresses used for subsequent transfers.

BWA BandWidth Allocation.

BWPBM BandWidth Portion Bit Map.

CONSTRAINED UNPREDICTABLE

Where an instruction can result in UNPREDICTABLE behavior, the Armv8 architecture specifies a narrow range of permitted behaviors. This range is the range of CONSTRAINED UNPREDICTABLE behavior. All implementations that are compliant with the architecture must follow the CONSTRAINED UNPREDICTABLE behavior.

Execution at Non-secure EL1 or EL0 of an instruction that is CONSTRAINED UNPREDICTABLE can be implemented as generating a trap exception that is taken to EL2, provided that at least one instruction that is not UNPREDICTABLE and is not CONSTRAINED UNPREDICTABLE causes a trap exception that is taken to EL2.

In body text, the term CONSTRAINED UNPREDICTABLE is shown in SMALL CAPITALS.

See also UNPREDICTABLE.

Core See *Processing element (PE)*.

CPBM Cache-Portion Bit Map.

CSU Cache-Storage Usage.

Downstream Information propagating in the direction from Requesters towards terminating Completer components.

DSB Data Synchronization Barrier.

E2H EL2 Host. A bit field in the HCR_EL2 register. This configuration executes a type-2 hypervisor and its host operating system in EL2 rather than EL1, for better performance.

Type-2 hypervisors run on a host operating system rather than running as a small, standalone OS-like program. For example, kvm is a type-2 hypervisor.

HCR An abbreviated reference to the Hypervisor Configuration Registers in AArch64 HCR_EL2 and in AArch32 HCR and HCR2.

ICN InterConnect Network.

ID An identifier or label.

Intermediate physical address (IPA)

An implementation of virtualization, the address to which a Guest OS maps a VA. A hypervisor might then map the IPA to a PA. Typically, the Guest OS is unaware of the translation from IPA to PA.

See also Physical address (PA), Virtual address (VA).

IPA See *Intermediate physical address (IPA)*.

kvm Kernel-based Virtual Machine, an open-source software package that implements a type-2 hypervisor within Linux.

LPI Locality-specific Peripheral Interrupt.

MBWU Memory BandWidth Usage.

Memory-system component

MSC. A function, unit, or design block in a memory system that can have partitionable resources. MSCs consist of all units that handle load or store requests issued by any MPAM Requester. These include cache memories, interconnects, memory management units, memory channel controllers, queues, buffers, rate adaptors, etc. An MSC may contain one or more resources that each may have zero or more resource partitioning controls. For example, a PE may contain several caches, each of which might have zero or more resource partitioning controls.

Memory-system resource

A resource that affects the performance of software's use of the memory system and is either local to an MSC (such as cache-memory capacity) or non-local (such as memory bandwidth, which is present over an entire path, from Requester to Completer, that may pass through multiple MSCs).

MMR	Memory-mapped Register.
MPAM	Memory system resource Partitioning and Monitoring.
MPAM information	The MPAM information bundle, comprising PARTID, PMG, and MPAM_NS.
MPAM_NS	MPAM security-space bit. It is not stored in a PE register; it comes from the current Security state of a PE and is communicated to MSCs as part of the MPAM information bundle. In non-PE Requesters, the Security state can be determined in other ways.
MPAM resource partition	See Resource partition .
MPAM_SP	In MPAM for RME the MPAM PARTID space indication.
MSC	Memory-system Component. See Memory-system component .
MSI	Message signaled interrupts. Signaled using a memory write that is usually directed at an interrupt translation service.
NRDY	Not-Ready bit. MPAM resource monitors set this bit to indicate that the monitor register does not currently have an accurate value.
NS	Non-Secure. A bit indicating that an address space is not Secure.
PA	See Physical address (PA) .
PARTID	The partition number component of an MPAM resource partition ID. See Resource partition
Partition	A division of resources. A partition is manifest in a PARTID and MPAM_NS. In an MSC, the PARTID and MPAM_NS select partitioning control settings that affect the partitioning by regulating the allocation of the resource to requests using that PARTID and MPAM_NS.
PE	See Processing element (PE) .
Physical address (PA)	An address that identifies a location in the physical memory map. See also Intermediate physical address (IPA) , Virtual address (VA) .
Physical PARTID	A partition ID that is transmitted with memory requests and can be used by MSCs to control resources usage. A physical PARTID is in either the Non-secure or Secure PARTID space. If MPAM for RME is implemented, there are two additional PARTID spaces, Realm PARTID space and Root PARTID space.
PMG	Performance Monitoring Group, a property of a partition used in MSCs by MPAM performance monitors that can be programmed to be sensitive to the particular PARTID and PMG combination.
Portion	A uniquely identifiable part of the resource. It is of fixed size or capacity. A particular resource has a constant number of portions. Portions are distinct. Portion n is the same part of the resource for every partition. Thus, every partition that is given access to a portion n shares access to portion n.
PPI	Private Peripheral Interrupt.
Processing element (PE)	The abstract machine defined in the Arm architecture, as documented in an Arm Architecture Reference Manual. A PE implementation compliant with the Arm architecture must conform with the behaviors described in the corresponding Arm Architecture Reference Manual.
RAZ	See Read-As-Zero (RAZ) .
RAZ/WI	Read-As-Zero, Writes Ignored. Hardware must implement the field as Read-as-Zero, and must ignore writes to the field. Software can rely on the field reading as all 0s, and on writes being ignored.

This description can apply to a single bit that reads as 0, or to a field that reads as all 0s.

See also *Read-As-Zero (RAZ)*.

Read-As-Zero (RAZ)

Hardware must implement the field as reading as all 0s.

Software:

- Can rely on the field reading as all 0s.
- Must use a SBZP policy to write to the field.

This description can apply to a single bit that reads as 0, or to a field that reads as all 0s.

See also *RAZ/WI, RES0*.

RES0

A reserved bit. Used for fields in register descriptions, and for fields in architecturally-defined data structures that are held in memory, for example in translation table descriptors.

Within the architecture, there are some cases where a register bit or field:

- Is RES0 in some defined architectural context.
- Has different defined behavior in a different architectural context.

———— **Note** —————

- RES0 is not used in descriptions of instruction encodings.
- Where an AArch32 System register is Architecturally mapped to an AArch64 System register, and a bit or field in that register is RES0 in one Execution state and has defined behavior in the other Execution state, this is an example of a bit or field with behavior that depends on the architectural context.

This means the definition of RES0 for fields in read/write registers is:

If a bit is RES0 in all contexts

For a bit in a read/write register, it is IMPLEMENTATION DEFINED whether:

1. The bit is hardwired to 0. In this case:
 - Reads of the bit always return 0.
 - Writes to the bit are ignored.
2. The bit can be written. In this case:
 - An indirect write to the register sets the bit to 0.
 - A read of the bit returns the last value successfully written, by either a direct or an indirect write, to the bit.
If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.
 - A direct write to the bit must update a storage location associated with the bit.
 - The value of the bit must have no effect on the operation of the PE, other than determining the value read back from the bit, unless this Manual explicitly defines additional properties for the bit.

Whether RES0 bits or fields follow behavior 1 or behavior 2 is IMPLEMENTATION DEFINED on a field-by-field basis.

If a bit is RES0 only in some contexts

For a bit in a read/write register, when the bit is described as RES0:

- An indirect write to the register sets the bit to 0.
- A read of the bit must return the value last successfully written to the bit, by either a direct or an indirect write, regardless of the use of the register when the bit was written.

If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.

- A direct write to the bit must update a storage location associated with the bit.
- While the use of the register is such that the bit is described as RES0, the value of the bit must have no effect on the operation of the PE, other than determining the value read back from that bit, unless this Manual explicitly defines additional properties for the bit.

Considering only contexts that apply to a particular implementation, if there is a context in which a bit is defined as RES0, another context in which the same bit is defined as RES1, and no context in which the bit is defined as a functional bit, then it is IMPLEMENTATION DEFINED whether:

- Writes to the bit are ignored, and reads of the bit return an UNKNOWN value.
- The value of the bit can be written, and a read returns the last value written to the bit.

The RES0 description can apply to bits or fields that are read-only, or are write-only:

- For a read-only bit, RES0 indicates that the bit reads as 0, but software must treat the bit as UNKNOWN.
- For a write-only bit, RES0 indicates that software must treat the bit as SBZ.

A bit that is RES0 in a context is reserved for possible future use in that context. To preserve forward compatibility, software:

- Must not rely on the bit reading as 0.
- Must use an SBZP policy to write to the bit.

This RES0 description can apply to a single bit, or to a field for which each bit of the field must be treated as RES0.

In body text, the term RES0 is shown in SMALL CAPITALS.

See also Read-As-Zero (RAZ), RES1, UNKNOWN.

RES1

A reserved bit. Used for fields in register descriptions, and for fields in architecturally-defined data structures that are held in memory, for example in translation table descriptors.

Within the architecture, there are some cases where a register bit or field:

- Is RES1 in some defined architectural context.
- Has different defined behavior in a different architectural context.

Note

- RES1 is not used in descriptions of instruction encodings.
- Where an AArch32 System register is Architecturally mapped to an AArch64 System register, and a bit or field in that register is RES1 in one Execution state and has defined behavior in the other Execution state, this is an example of a bit or field with behavior that depends on the architectural context.

This means the definition of RES1 for fields in read/write registers is:

If a bit is RES1 in all contexts

For a bit in a read/write register, it is IMPLEMENTATION DEFINED whether:

1. The bit is hardwired to 1. In this case:
 - Reads of the bit always return 1.
 - Writes to the bit are ignored.
2. The bit can be written. In this case:
 - An indirect write to the register sets the bit to 1.
 - A read of the bit returns the last value successfully written, by either a direct or an indirect write, to the bit.
If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.
 - A direct write to the bit must update a storage location associated with the bit.
 - The value of the bit must have no effect on the operation of the PE, other than determining the value read back from the bit, unless this Manual explicitly defines additional properties for the bit.

Whether RES1 bits or fields follow behavior 1 or behavior 2 is IMPLEMENTATION DEFINED on a field-by-field basis.

If a bit is RES1 only in some contexts

For a bit in a read/write register, when the bit is described as RES1:

- An indirect write to the register sets the bit to 1.
- A read of the bit must return the value last successfully written to the bit, regardless of the use of the register when the bit was written.

———— **Note** —————

As indicated in this list, this value might be written by an indirect write to the register.

—————
If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.

- A direct write to the bit must update a storage location associated with the bit.
- While the use of the register is such that the bit is described as RES1, the value of the bit must have no effect on the operation of the PE, other than determining the value read back from that bit, unless this Manual explicitly defines additional properties for the bit.

Considering only contexts that apply to a particular implementation, if there is a context in which a bit is defined as RES0, another context in which the same bit is defined as RES1, and no context in which the bit is defined as a functional bit, then it is IMPLEMENTATION DEFINED whether:

- Writes to the bit are ignored, and reads of the bit return an UNKNOWN value.
- The value of the bit can be written, and a read returns the last value written to the bit.

The RES1 description can apply to bits or fields that are read-only, or are write-only:

- For a read-only bit, RES1 indicates that the bit reads as 1, but software must treat the bit as UNKNOWN.
- For a write-only bit, RES1 indicates that software must treat the bit as SBO.

A bit that is RES1 in a context is reserved for possible future use in that context. To preserve forward compatibility, software:

- Must not rely on the bit reading as 1.
- Must use an SBOP policy to write to the bit.

This RES1 description can apply to a single bit, or to a field for which each bit of the field must be treated as RES1.

In body text, the term RES1 is shown in SMALL CAPITALS.

See also RES0, UNKNOWN.

Reserved

Unless otherwise stated:

- Instructions that are reserved or that access reserved registers have UNPREDICTABLE or CONSTRAINED UNPREDICTABLE behavior.
- Bit positions described as reserved are:
 - In an RW or WO register, RES0.
 - In an RO register, UNK.

Resource partition

The collection of MPAM resource control settings associated with a software environment and identified by the combination of a physical PARTID space and a partition number.

RIS

Resource instance selection. The value in `MPAMCFG_PART_SEL.RIS` selects the resource instance that is configured through `MPAMCFG_*` registers and described by the MPAMF ID registers. See *RIS controls in MPAMCFG_PART_SEL*.

RME

Realm Management Extension. RME specifies how PE execution context is mapped to Security states.

SCR

Part of the name of a Secure Configuration Register.

SMMU

System Memory-Management Unit.

SPE

Statistical Profiling Extension.

SPI

Shared Peripheral Interrupt.

TGE

Trap General Exception. A field in the `HCR_EL2` register. It causes EL0 exceptions, that would normally trap to EL1, to instead trap to EL2. This function can be used to run an EL2 host's applications at EL0, so that any exceptions in the application trap to the host OS at EL2.

UNDEFINED

Indicates cases where an attempt to execute a particular encoding bit pattern generates an exception, that is taken to the current Exception level, or to the default Exception level for taking exceptions if the UNDEFINED encoding was executed at EL0. This applies to:

- Any encoding that is not allocated to any instruction.
- Any encoding that is defined as never accessible at the current Exception level.
- Some cases where an enable, disable, or trap control means an encoding is not accessible at the current Exception level.

If the generated exception is taken to an Exception level that is using AArch32 then it is taken as an Undefined Instruction exception.

————— Note —————

On reset, the default Exception level for taking exceptions from EL0 is EL1. However, an implementation might include controls that can change this, effectively making EL1 inactive. See the description of the Exception model for more information.

In body text, the term UNDEFINED is shown in SMALL CAPITALS.

UNKNOWN	<p>An UNKNOWN value does not contain valid data, and can vary from moment to moment, instruction to instruction, and implementation to implementation. An UNKNOWN value must not return information that cannot be accessed at the current or a lower level of privilege using instructions that are not UNPREDICTABLE, are not CONSTRAINED UNPREDICTABLE, and do not return UNKNOWN values.</p> <p>An UNKNOWN value must not be documented or promoted as having a defined value or effect.</p> <p>In body text, the term UNKNOWN is shown in SMALL CAPITALS.</p> <p><i>See also</i> CONSTRAINED UNPREDICTABLE, UNDEFINED, UNPREDICTABLE.</p>
UNPREDICTABLE	<p>Means the behavior cannot be relied upon. UNPREDICTABLE behavior must not perform any function that cannot be performed at the current or a lower level of privilege using instructions that are not UNPREDICTABLE.</p> <p>UNPREDICTABLE behavior must not be documented or promoted as having a defined effect.</p> <p>An instruction that is UNPREDICTABLE can be implemented as UNDEFINED.</p> <p>Execution at Non-secure EL1 or EL0 of an instruction that is UNPREDICTABLE can be implemented as generating a trap exception that is taken to EL2, provided that at least one instruction that is not UNPREDICTABLE and is not CONSTRAINED UNPREDICTABLE causes a trap exception that is taken to EL2.</p> <p>In body text, the term UNPREDICTABLE is shown in SMALL CAPITALS.</p> <p><i>See also</i> CONSTRAINED UNPREDICTABLE, UNDEFINED.</p>
Upstream	Information propagating in the direction from terminating Completer components towards Requesters.
VA	<i>See</i> Virtual address (VA) .
Virtual address (VA)	<p>An address generated by an Arm PE. This means it is an address that might be held in the program counter of the PE. For a PMSA implementation, the virtual address is identical to the physical address.</p> <p><i>See also</i> Intermediate physical address (IPA), Physical address (PA).</p>
Virtual PARTID	One of a small range of PARTIDs that can be used by a virtual machine (VM). Virtual PARTIDs are mapped into physical PARTIDs using the virtual partition mapping entries in the MPAMVPM0 - MPAMVPM7 registers.
VM	Virtual Machine.
VMM	Virtual Machine Monitor. An alias for “hypervisor”.
Word	A 32-bit data item. Words are normally word-aligned in Arm systems.
Word-aligned	Means that the address is divisible by 4.