

OpenPOWER Memory Bus Specification for OMI

Memory Function Unit Interface

Workgroup Specification

Revision 1.93 (April 26, 2021)

REVIEW DRAFT



www.openpowerfoundation.org

OpenPOWER Memory Bus Specification for OMI: Memory Function Unit Interface

Memory Work Group <memwg-chair@openpowerfoundation.org>
OpenPower Foundation

Revision 1.93 (April 26, 2021)

Copyright © 2019, 2020 OpenPOWER Foundation

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

The limited permissions granted above are perpetual and will not be revoked by OpenPOWER or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE OpenPOWER Foundation AS WELL AS THE AUTHORS AND DEVELOPERS OF THIS STANDARDS FINAL DELIVERABLE OR OTHER DOCUMENT HEREBY DISCLAIM ALL OTHER WARRANTIES AND CONDITIONS, EITHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES, DUTIES OR CONDITIONS OF MERCHANTABILITY, OF FITNESS FOR A PARTICULAR PURPOSE, OF ACCURACY OR COMPLETENESS OF RESPONSES, OF RESULTS, OF WORKMANLIKE EFFORT, OF LACK OF VIRUSES, OF LACK OF NEGLIGENCE OR NON-INFRINGEMENT.

OpenPOWER, the OpenPOWER logo, and openpowerfoundation.org are trademarks or registered trademarks of OpenPOWER Foundation, Inc., registered in many jurisdictions worldwide. Other company, product, and service names may be trademarks or service marks of others.

Abstract

This document is a Standard Track, Work Group Specification work product owned by the Memory Workgroup and handled in compliance with the requirements outlined in the *OpenPOWER Foundation Work Group (WG) Process* document. It was created using the *Master Template Guide* version 1.0.0. Comments, questions, etc. can be submitted to the public mailing list for this document at <memwg-opmb@mailinglist.openpowerfoundation.org>.

The OpenPOWER Memory Bus Specification defines the OpenPOWER Memory Bus Architecture for the development of Memory Function Units(MFU) and the integration of those MFUs into the OpenPOWER system structure. An MFU is a logic block developed by a member of the OpenPOWER eco-system to enable the integration of specialized memory technology and functional processing for data stored in the memory technology.

The OpenPOWER Memory Bus Specification is a standards track work product of the OpenPOWER Memory Workgroup.

List of Figures

2.1. An Abstract OPMB-Compliant Memory System	5
3.1. OPMB Downstream Signal Interface	9
3.2. OPMB Upstream Signal Interface	9
3.3. OPMB Inband Configuration Signal Interface	10
4.1. I2C Boot Config Command	20
4.2. I2C Boot Config Command Encodings	21
4.3. I2C Boot Config Command Status Encodings	21
4.4. I2C Status Command	22
4.5. I2C Status Command Response Encodings	22
4.6. I2C Register Address Latch Command	23
4.7. I2C Register Address Latch Command Status Encodings	23
4.8. I2C Register Read Command	24
4.9. I2C Read Command Status Encodings	24
4.10. I2C Register Write Command	25
4.11. I2C Write Command Status Encodings	25
6.1. OMI Configuration Register Initialization Flow	41
8.1. High Level Firmware Initialization Flow	44

List of Tables

1.1. Register References	3
2.1. Sizes of Main Storage Address Spaces	6
3.1. MSL Interface Variations	8
3.2. Downstream MSL Command and Data Signals	10
3.3. Upstream MSL Command and Data Signals	11
3.4. MSL Inband Configuration Signals	13
4.1. OpenCAPI 3.1 Device Interface Class Commands	16
6.1. Required registers	27
6.2. OpenCAPI Function 0 Device ID & Vendor ID Register	28
6.3. OpenCAPI Function 0 MMIO Subsystem ID & Subsystem Vendor ID Register	28
6.4. OpenCAPI Function 1 MMIO Base Upper Address 0 Register	28
6.5. OpenCAPI Function 1 MMIO Base Lower Address 0 Register	29
6.6. OpenCAPI Function 0 Device ID & Vendor ID Register	29
6.7. OpenCAPI Function 1 Memory Space Enable Register	29
6.8. OpenCAPI Function 1 MMIO Subsystem ID & Subsystem Vendor ID Register	30
6.9. OpenCAPI AFU Control DVSEC AFU Enable and ID Register	30
6.10. OpenCAPI AFU Control DVSEC acTAG Register	30
6.11. OpenCAPI AFU Control DVSEC acTAG Length Register	31
6.12. OpenCAPI Function acTAG Register	31
6.13. OpenCAPI AFU Control DVSEC PASID Base Register	31
6.14. OpenCAPI AFU Control DVSEC PASID Length Register	32
6.15. OpenCAPI TLX Receive Template 31-0 Capabilities Register	32
6.16. OpenCAPI TLX Flit Receive Rate per Template 7-0 Capabilities Register	33
6.17. OpenCAPI TLX Transmit Template Configurations Register	33
6.18. OpenCAPI TLX Flit Receive Rate per Template 7-0 Capabilities Register	35
6.19. OpenCAPI Transport Layer Configuration Register	35
6.20. ICE Configuration Register 0	36
6.21. ICE Configuration Register 1	36
6.22. ICE Configuration Register 2	37
6.23. ICE Interrupt Handle Register 0	38
6.24. ICE Interrupt Handle Register 1	38
6.25. ICE Interrupt Handle Register 2	38
6.26. ICE Interrupt Handle Register 3	38
6.27. ICE Trap Register 0	39
6.28. ICE Trap Register 2	39
6.29. ICE Trap Register 3	40
6.30. ICE Trap Register 5	40
6.31. ICE Trap Register 6	40
6.32. ICE Trap Register 8	41
7.1. OpenPOWER System SUE Meta Bit Encodings	42

Document links

Document links frequently appear throughout the documents. Generally, these links include a text for the link, followed by a page number in parenthesis. For example, this link, [Preface \[vi\]](#), references the [Preface](#) chapter on page [vi](#).

2. Document change history

This version of the guide replaces and obsoletes all earlier versions.

The following table describes the most recent changes:

Revision Date	Summary of Changes
October 11, 2020	<ul style="list-style-type: none">Revision 0.91 - Draft to review for voting

2. Field names and variables are written in *italic type*. Required parameters are enclosed in angle brackets. Optional parameters are enclosed in brackets. For example: *afu* <*f*,*b*>_*wr*[*a*].

This document uses the following symbols:

&	Bitwise AND
	Bitwise OR
~	Bitwise NOT
%	Modulus
=	Equal to
!=	Not equal to
>=	Greater than or equal to
≤	Less than or equal to
x >> y	Shift to the right; for example, 6 >> 2 = 1; least-significant y bits are dropped
x << y	Shift to the left; for example, 3 << 2 = 12; least-significant y bits are replaced zeros
	Concatenate

1.5. References to Registers, Fields, and Bits

Registers are referred to by their full name or by their short name (also called the register mnemonic). Fields are referred to by their field name or by their bit position. The following table describes how registers, fields, and bit ranges are referred to in this document and provides examples.

Table 1.1. Register References

Type of Reference	Format	Example
Reference to a specific register and a specific field using the register short name and the field name	Register_Short_Name[Field_Name]	MSR[R]
Reference to a field using the field name	[Field_Name]	[R]
Reference to a specific register and to multiple fields using the register short name and the field names	Register_Short_Name[Field_Name1, Field_Name2]	MSR[FE0, FE1]
Reference to a specific register and to multiple fields using the register short name and the bit positions.	Register_Short_Name[Bit_Number, Bit_Number]	MSR[52, 55]
Reference to a specific register and to a field using the register short name and the bit position or the bit range.	Register_Short_Name[Bit_Number]	MSR[52]
	Register_Short_Name[Starting_Bit_Number:Ending_Bit_Number]	MSR[39:44]
A field name followed by an equal sign (=) and a value indicates the value for that field.	Register_Short_Name[Field_Name]= <i>n</i>	MSR[FE0] = '1'
	Register_Short_Name[Bit_Number]= <i>n</i>	MSR[FE] = x'1'
		MSR[52] = '0'
	Register_Short_Name[Starting_Bit_Number:Ending_Bit_Number]= <i>n</i>	MSR[52] = x'0'
		MSR[39:43] = '10010'
		MSR[39:43] = x'11'
Where <i>n</i> is the binary or hexadecimal value for the field or bits specified in the brackets.		

2. Introduction

The OpenPOWER Memory Bus (OPMB) Specification for OMI (Open Memory Interface based on OpenCAPI 3.1) defines a memory function unit interface structure for attaching one or more memory function units (MFUs) to an OpenPOWER™ System with OpenCAPI using a Open Memory Interface (OMI). The intent for the creation of this documentation is to foster the development of an ecosystem and enable the implementation of different memory buffering devices that support different MFUs and enable the use of different memory technologies in different configurations.

Physically, an OPMB-compliant memory buffering device can consist of a single chip, a multi-chip module (or modules), or multiple single-chip modules on a system board or other second-level package. The design depends on the technology used, and on the cost and performance characteristics of the intended design point.

Figure 2.1. An Abstract OPMB-Compliant Memory System

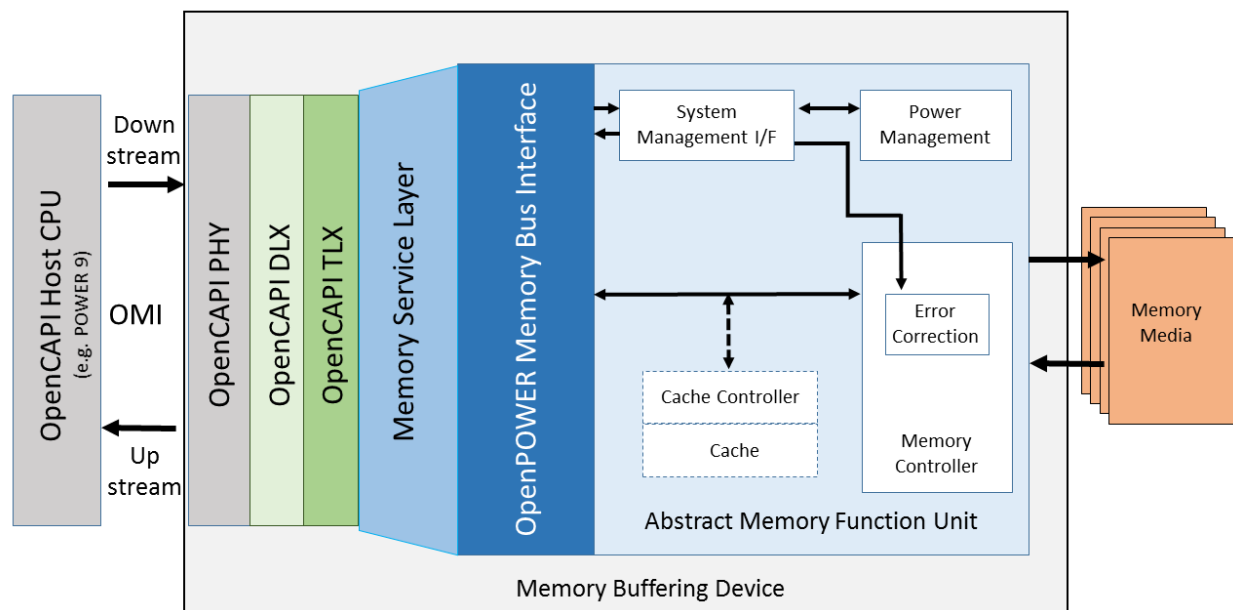


Figure 2.1, “An Abstract OPMB-Compliant Memory System” [5] illustrates a OPMB-compliant memory system using OpenCAPI, including: a host CPU with a Open Memory Interface (OMI), a memory buffering device that receives commands and data as well as return data and status information through the OMI to the host CPU, and memory media controlled by the memory buffering device.

2.1. Open Memory Interface (OMI)

The Open Memory Interface (OMI) is a physical layer interface that enables the transport of memory command, address and data encapsulated in flits through high speed Serdes links to and from the host CPU. The OMI is an open standard memory interface implemented using either the OpenCAPI

Figure 3.1. OPMB Downstream Signal Interface

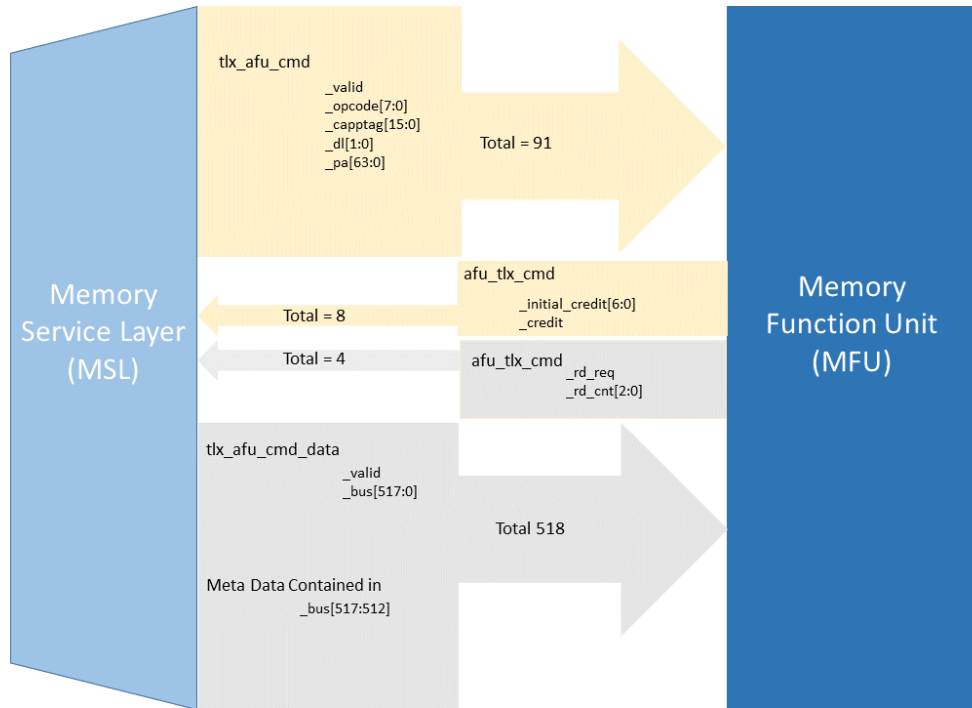


Figure 3.1, "OPMB Downstream Signal Interface" [9] Illustration of the OPMB downstream signal interface between the MSL and MFU

Figure 3.2. OPMB Upstream Signal Interface

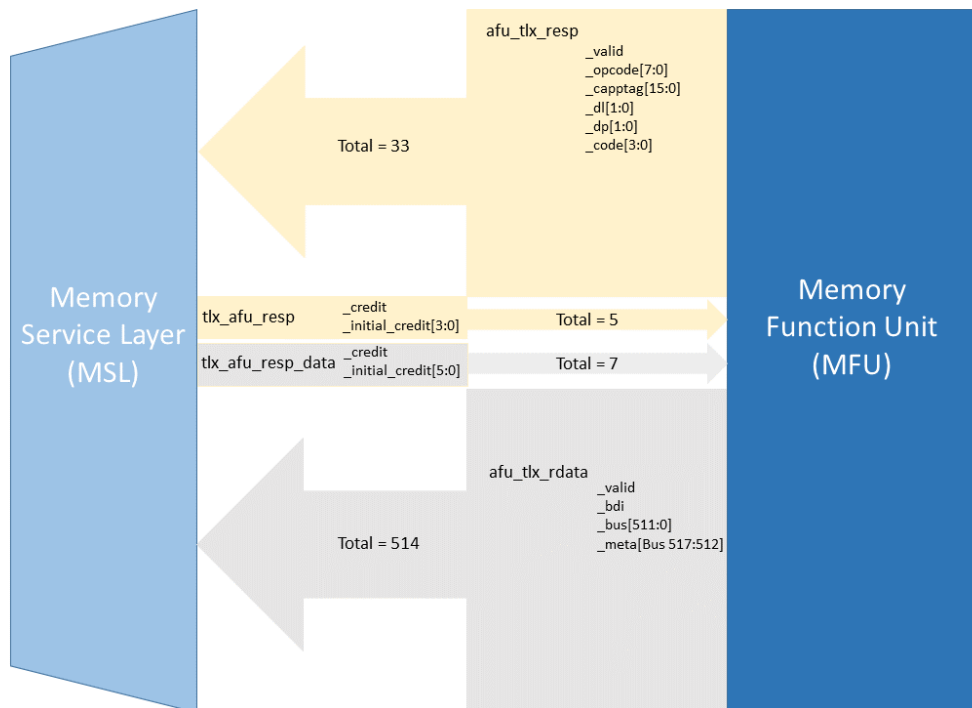


Figure 4.2. I2C Boot Config Command Encodings

- 31:16-> Reserved for future use
- 14:13 -> Boot Mode
 - 0b00 = Normal Boot
 - 0b01 Manufacturing Mode
 - 0b10 Product Qual Mode
 - 0b11 Reserved
- 12 -> OMI Loopback Testing
 - 0b0 = No loopback testing
 - 0b1 = Perform loopback testing
- 11:10 -> Transport Layer
 - 0b00 = OMI
 - All others reserved
- 9:8 -> Boot Config Command Stage
 - 0b00 = Stage 0: Only SerDes will be Initialized
 - 0b01 = Stage 1: DL Training will be run
 - All others reserved
- 7 -> DFE Enable if supported
- 6:4 -> Lane Mode
 - 0b000 = Undefined, error
 - 0b001 = 8 Lane
 - 0b010 = 4 Lane
 - All others reserved
- 3:0 -> OMI Frequency
 - 0b0000 = Undefined, error
 - 0b0001 = 21.33 Gbps
 - 0b0010 = 23.46 Gbps
 - 0b0011 = 25.6 Gbps
 - All others reserved

After issuing a Boot Config command the status shall be checked by issuing a status command over I2C. The following figure shows the status bit encodings for the boot config command.

Figure 4.3. I2C Boot Config Command Status Encodings

- 0x00 = Success
- 0x01 = Invalid Boot Config Command
- 0x02 = Loopback Test Fail
- 0x03 = OMI SerDes Init Fail
- 0x04 = DLx Config Fail
- 0x10 = Lane Inversion configuration fail
- 0x20 = SerDes hardware parity and UECC errors
- 0xFE = I2C Interface Busy

4.2.2. I2C Status Command

The Status Command is an integral part of the initialization process. After deasserting reset to the OMI channel the system begins to poll the memory buffer using the status command. Once status indicates a runtime mode the system will continue with initialization. In order to issue an I2C Status Command use the following flow on the I2C bus.

Figure 4.8. I2C Register Read Command

- Address Latch Command
 - START
 - 7b Device Address + 1b R/W# bit
 - Memory Buffer 0b010,0000 + R/W# 0b0 = 0x40
 - Byte 0 = Command ID = 0x03 (Address Latch Command)
 - Byte 1 = Number of bytes to follow this byte = 0x04 for a 32b address
 - Byte 2 = Byte 0 of 32b address -> Left most byte of 32b address
 - Byte 3 = Byte 1 of 32b address
 - Byte 4 = Byte 2 of 32b address
 - Byte 5 = Byte 3 of 32b address -> Right most byte of 32b address
 - STOP
- Read Command
 - START
 - 7b Device Address + 1b R/W# bit
 - Memory Buffer 0b010,0000 + R/W# 0b0 = 0x40
 - Byte 0 = Command ID = 0x04 (Read Command)
 - Byte 1 = Number of bytes to follow this byte = 0x04 for a 32b address
 - Byte 2 = Byte 0 of 32b address -> Left most byte of 32b address
 - Byte 3 = Byte 1 of 32b address
 - Byte 4 = Byte 2 of 32b address
 - Byte 5 = Byte 3 of 32b address -> Right most byte of 32b address
 - START
 - 7b Device Address + 1b R/W# bit
 - Memory Buffer 0b010,0000 + R/W# 0b1 = 0x41
 - Byte 0 = Length -> should receive a 0x04
 - Byte 1 = Data byte 0
 - Byte 2 = Data byte 1
 - Byte 3 = Data byte 2
 - Byte 4 = Data byte 3
 - STOP

After issuing a Read command the status should be checked by issuing a status command over I2C. The following figure shows the status bit encodings for the Read command.

Figure 4.9. I2C Read Command Status Encodings

0x00 = Success
<address> = Invalid address

4.2.5. I2C Register Write Command

In order to issue an I2C Register Write refer to the following flow on the I2C bus.

8b command ID encoding = 0x05

Figure 4.10. I2C Register Write Command

- Write Command
 - START
 - 7b Device Address + 1b R/W# bit
 - Memory Buffer 0b010,0000 + R/W# 0b0 = 0x40
 - Byte 0 = Command ID = 0x05 (Write Command)
 - Byte 1 = Number of bytes to follow this byte = 0x08 for a 32b address + 4 bytes data
 - Byte 2 = Byte 0 of 32b address -> Left most byte of 32b address
 - Byte 3 = Byte 1 of 32b address
 - Byte 4 = Byte 2 of 32b address
 - Byte 5 = Byte 3 of 32b address -> Right most byte of 32b address
 - Byte 6 = Data byte 0
 - Byte 7 = Data byte 1
 - Byte 8 = Data byte 2
 - Byte 9 = Data byte 3
 - STOP

After issuing a Write command the status should be checked by issuing a status command over I2C. The following figure shows the status bit encodings for the Write command.

Figure 4.11. I2C Write Command Status Encodings

0x00 = Success
0xC0 = Error

0	0x0	ICETCFG_TEMPLATE_0	RW	When set indicates that Template 0 is supported. This field is disruptive to write during runtime.
---	-----	--------------------	----	--

Table 6.18. OpenCAPI TLX Flit Receive Rate per Template 7-0 Capabilities Register

RTL Mnemonic		AFU_MAC.CFG.ICETRATE		
Name		OpenCAPI TLX Flit Transmit Rate per Templates 7-0 Configurations Register		
Reset Value		0x00000000		
Description		Indicates the transmit rate that can be supported for templates 7-0 for the design		
Address Type		Address		
Configuration		0x0000026c		
Bit Range	Initial Value	Field Name	Field Type	Description
31:28	0x0	ICERRATE_TEMPLATE_7	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 7. A value of zero indicates that a control flit can be sent in the following flit cycle. This field can be freely written during runtime.
27:24	0x0	ICETRATE_TEMPLATE_6	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 6. A value of zero indicates that a control flit can be sent in the following flit cycle. This field can be freely written during runtime.
23:20	0x0	ICETRATE_TEMPLATE_5	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 5. A value of zero indicates that a control flit can be sent in the following flit cycle. This field is disruptive to write during runtime.
19:16	0x0	ICETRATE_TEMPLATE_4	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 4. A value of zero indicates that a control flit can be sent in the following flit cycle. This field can be freely written during runtime.
15:12	0x0	ICETRATE_TEMPLATE_3	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 3. A value of zero indicates that a control flit can be sent in the following flit cycle. This field can be freely written during runtime.
11:8	0x0	ICERRATE_TEMPLATE_2	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 2. A value of zero indicates that a control flit can be sent in the following flit cycle. This field can be freely written during runtime.
7:4	0x0	ICERRATE_TEMPLATE_1	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 1. A value of zero indicates that a control flit can be sent in the following flit cycle. This field is disruptive to write during runtime.
3:0	0x0	ICERRATE_TEMPLATE_0	RW	Indicates the size of gap (measured in flits) before the next control flit can be sent for Template 0. A value of zero indicates that a control flit can be sent in the following flit cycle. This field is disruptive to write during runtime.

Table 6.19. OpenCAPI Transport Layer Configuration Register

RTL Mnemonic	AFU_MAC.CFG.ICEVER
Name	OpenCAPI Transport Layer Configuration Register

2	0x0	ICECFG1_INJECT_MCHK	RW	When set will inject aapplication interrupt to be reported to the host via the interrupt interface Interrupts are generated on the rising edge of an error. To inject a second interrupt of this type the bit will need to be cleared and then set again.
3	0x0	ICECFG1_TRAP_UPDATE	RW	when set, trap updates every cycle and does not freeze on error
4	0x0	ICECFG1_HOLD_ACUM	RW	when set, hold reg acumulates on error and does not freeze on fir error out
5:15	0x0	Reserved	RO	Constant 0b00000000000
16:20	0x0	ICECFG1_EDPL_RESERVED0	RW	EDPL Reserved, Constant 0b0000
21	0x1	ICECFG1_EDPL_ENABLE	RW	EDPL Enable
22	0x0	ICECFG1_EDPL_MAX_COUNTER_RESET	RW	EDPL Max Counter Reset
23:31	0x0	ICECFG1_EDPL_RESERVED1	RW	EDPL reserved
32:39	0x0	ICECFG1_EDPL_INJECT_PER_LANE	RW	EDPL Inject Per lane
40	0x0	ICECFG1_EDPL_RESERVED2	RW	EDPL reserved
41:43	0x4	ICECFG1_EDPL_ERROR_THRESHOLD	RW	EDPL Error Threshold
44:47	0x5	ICECFG1_EDPL_TIME_WINDOW	RW	EDPL Time Window
48:51	0x0	ICECFG1_INTRP_CMDFLAG_0	RW	ICE Interrupt handle 0 command flag. Fills the cmd_flag portion of the intrp_re_q command found in section 2.3 of the OpenCAPI TL 3.1 Specification. Used for Channel Checkstop Interrupt Request.
52:55	0x0	ICECFG1_INTRP_CMDFLAG_1	RW	ICE Interrupt handle 1 command flag Fills the 4-bit cmd_flag portion of the intrp_req command found in section 2.3 of the OpenCAPI TL 3.1 Specification. Used for Recoverable Attention Interrupt Request.
56:59	0x0	ICECFG1_INTRP_CMDFLAG_2	RW	ICE Interrupt handle 2 command flag Fills the cmd_flag portion of the intrp_req command found in section 2.3 of the OpenCAPI TL 3.1 Specification. Used for Special Attention Interrupt Request.
60:63	0x0	ICECFG1_INTRP_CMDFLAG_3	RW	ICE Interrupt handle 3 command flag Fills the cmd_flag portion of the intrp_req command found in section 2.3 of the OpenCAPI TL 3.1 Specification. Used for Application Interrupt Request.

Table 6.22. ICE Configuration Register 2

RTL Mnemonic		AFU_MAC.MMIO.ICE_REGS.ICECFG2		
Name		ICE Configuration Register 2		
Reset Value		0x0000000000000000		
Description		NA		
Address Type		Address		
Configuration		0x0801240e		
Bit Range	Initial Value	Field Name	Field Type	Description
0:31	0x0	ICECFG2_ECID	RO	Unique identifier for this design.
32:63	0x0	ICECFG2_DEVICE_ID_REG	RO	Device Id information for the design. For example... 32:35 = Major Release 36:39 = Minor Release 44:47 = Location Code 56:63 = Chip ID

Configuration		0x08012426		
Bit Range	Initial Value	Field Name	Field Type	Description
0:63	0x0	ICETRAP2_DEBUG_BUS	ROX	MMIO/error Reporting Trace 0 = xlate_mmio_cmd_val 1 = xlate_mmi_cmd_rd 2:17 = xlate_mmio_cmd_tag 18 = cmd_pop_q 19 = cmd_data_out(0) 20:35 = cmd_data_out(51:36) 36:53 = sc_addr_v 54:56 = cmd_addr_offset(2:0) 57:58 = intrp_req_chan_xstop 59:60 = intrp_req_app_intrp

Table 6.29. ICE Trap Register 3

RTL Mnemonic		AFU_MAC.MMIO.ICE_REGS.ICETRAP3		
Name		ICE Trap Register 3		
Reset Value		0x0000000000000000		
Description		NA		
Address Type		Address		
Configuration		0x08012427		
Bit Range	Initial Value	Field Name	Field Type	Description
0:63	0x0	ICETRAP3_DEBUG_BUS	ROX	TLX Framer Trace 0 = rsp_valid_d1 1 = injct_cfg_rsp 2:4 = rsp_packet_reg(2:0) 5:20 = rsp_packet_reg(23:8) 21 = valid_rsp_leaving_fifo 22:24 = vc0_fifo_output_reg(2:0) 25:40 = vc0_fifo_output_reg(23:8) 41:43 = rsp_tmpl 44 = rsp_cf_was_taken 45:48 = rcf_df_count_last 49:52 = meta_run_length 53:56 = framer_run_length 57 = metadata_enabled 58:60 = meta_packet_ptr(2:0) 61:63 = meta_packet_ptr(2:0)

Table 6.30. ICE Trap Register 5

RTL Mnemonic		AFU_MAC.MMIO.ICE_REGS.ICETRAP5		
Name		ICE Trap Register 5		
Reset Value		0x0000000000000000		
Description		NA		
Address Type		Address		
Configuration		0x08012429		
Bit Range	Initial Value	Field Name	Field Type	Description
0:31	0x0	ICETRAP5_DLX_DEBUG_REG04	ROX	DLX Debug Reg 04
32:63	0x0	ICETRAP5_DLX_DEBUG_REG05	ROX	DLX debug Reg 05

Table 6.31. ICE Trap Register 6

RTL Mnemonic		AFU_MAC.MMIO.ICE_REGS.ICETRAP6		
Name		ICE Trap Register 6		
Reset Value		0x0000000000000000		
Description		NA		
Address Type		Address		
Configuration		0x0801242a		
Bit Range	Initial Value	Field Name	Field Type	Description

0:31	0x0	ICETRAP6_DLX_DEBUG_REG06	ROX	DLX Debug Reg06
32:63	0x0	ICETRAP6_DLX_DEBUG_REG07	ROX	DLX Debug Reg07

Table 6.32. ICE Trap Register 8

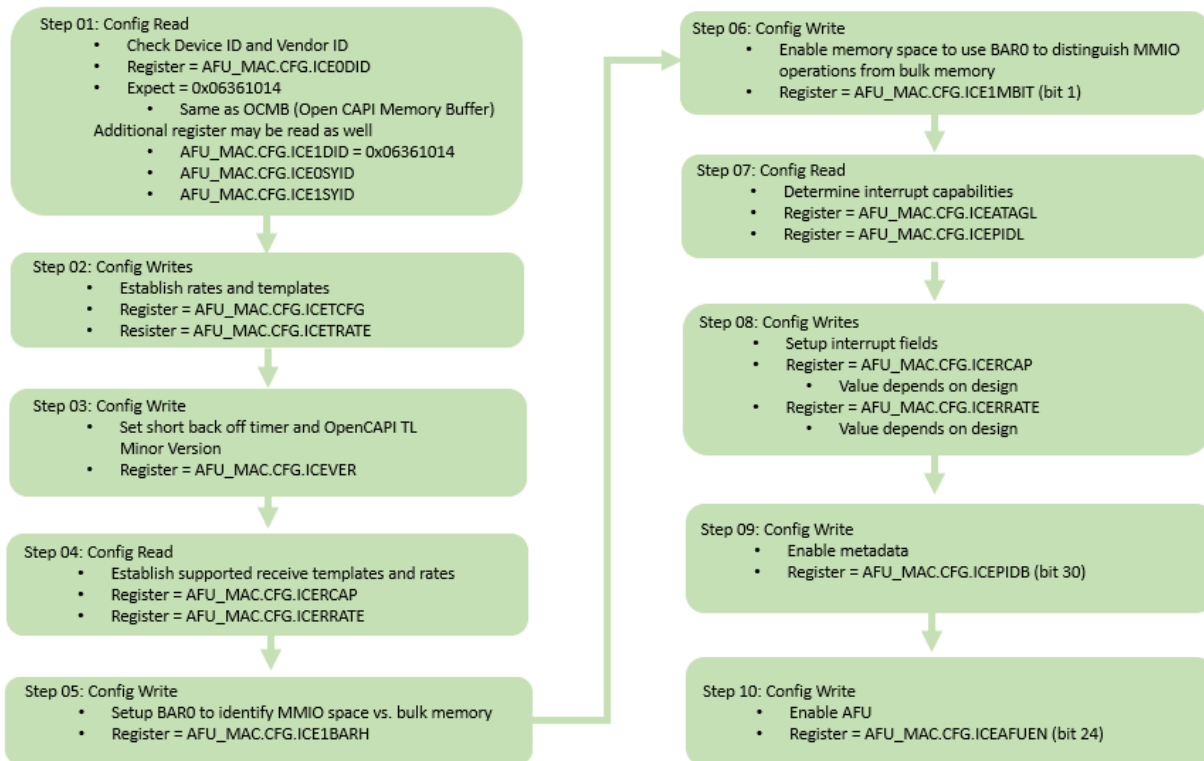
RTL Mnemonic		AFU_MAC.MMIO.ICE_REGS.ICETRAP8		
Name		ICE Trap Register 8		
Reset Value		0x0000000000000000		
Description		NA		
Address Type		Address		
Configuration		0x0801242c		
Bit Range	Initial Value	Field Name	Field Type	Description
0:63	0x0	ICETRAP8_DEBUG_BUS	ROX	The debug bus register fields and definition are MFU dependent and outside the scope of this document. Refer to the example design Genini/ICE for example guidance.

6.3. Configuration Register Initialization

After the OMI link is trained and the checked for errors, the OMI configuration is performed. To understand where this step occurs in the initialization flow, refer to the initialization flow in chapter eight.

The following figure depicts the configuration register initialization flow.

Figure 6.1. OMI Configuration Register Initialization Flow



7. Meta Bit Requirements

This chapter describes the minimum metabit requirements when deploying an MFU connected to the OpenPOWER Memory Bus (OPMB) as attached to an OpenCAPI enabled OpenPower system. Some systems may require usage of additional metabits above the minimum stated here. Therefore, when designin an MFU, it is advisable to refer to the appropriate system level documentation for correct implementation of all required meta bits.

At a minimum the MFU will encode the metabits to distinguish between a failure on a downstream write transfer from the host or a failure on the read path from the attached media.

Encoding metabits for a read or write failure is called a SUE encoding. Two of the six metabits are deployed for SUE encoding, metabit 2 and metabit 5.

Downstream write failures shall be encoded as a write failure SUE anywhere in the path where a failure occurs starting with the data parity check on the downstream interface from the MSL to the MFU at a minimum. The MFU must also maintain an encoded SUE received in the metabits from the host on the MSL on the downstream MSL interface.

Upstream read failures shall be encoded as a media read failure SUE anywhere in the path where a failure occurs ending with a data integrity check (ECC, parity, etc.) on the read data prior to transfer from the MFU to the MSL. Write path SUE must be maintained with the data and should be present on the read transaction.

7.1. Meta Bit Definitions

This section defines the meta bit encodings.

Table 7.1. OpenPOWER System SUE Meta Bit Encodings

Condition	Write Path Meta(2,5)	Read No-Media Error Meta(2,5)	Read Media Error Meta(2,5)	Description
Good data from MSL Meta Values, both 0	(0,0)	(0,0)	(0,1)	Good data from the host should always have matching meta bit values. The meta bit values are maintained if there is no read error from the media thus returning the same values back to the host. On a read fail from the media the meta bits are always set to Meta(2,5) = (0,1).
Good data from MSL Meta Values, both 1	(1,1)	(1,1)	(0,1)	Good data from the host should always have matching meta bit values. The meta bit values are maintained if there is no read error from the media thus returning the same values back to the host. On a read fail from the media the meta bits are always set to Meta(2,5) = (0,1)
Known-bad write data from MSL	(0,1) or (1,0)	(1,0)	(0,1)	Known bad data from the host has mismtaching meta bit values. These values will be stored in the media. When the data is retrieved from the media and the meta bits indicate stored bad data the meta bits must always be Meta(2,5)=1,0 even if they were stored as 0,1. If the data is read from the media and there is a fail on the read the meta bits will be encoded as Meta(2,5)=0,1. The reasoning for this is to differentiate if the data was corrupted in the write path from host to the media or during a read from the media to the host. If both the data

				was corrupted on the write and a read back of the poisoned data was erroneous then the read fail indication takes priority.
Uncorrectable error on write	Any value UE Decoded	(1,0)	(0,1)	Similar to when known bad data from the host is presented with mismatching meta bit values. In this case the write data from the TLX has bad parity or uncorrectable ECC encoding. In addition, data could be corrupted somewhere in the MFU prior to writing out to the media. In this case the MFU should mark the data with META(2,5)=(1,0) prior to writing to the media. When the data is retrieved from the media and the meta bits indicate stored bad data the meta bits must always be Meta(2,5)=1,0. If the data is read from the media and there is a fail on the read the meta bits will be encoded as Meta(2,5)=0,1. The reasoning for this is to differentiate if the data was corrupted in the write path from host to the media or during a read from the media to the host. If both the data was corrupted on the write and a read back of the poisoned data was erroneous then the read fail indication takes priority.

8. Firmware Interaction

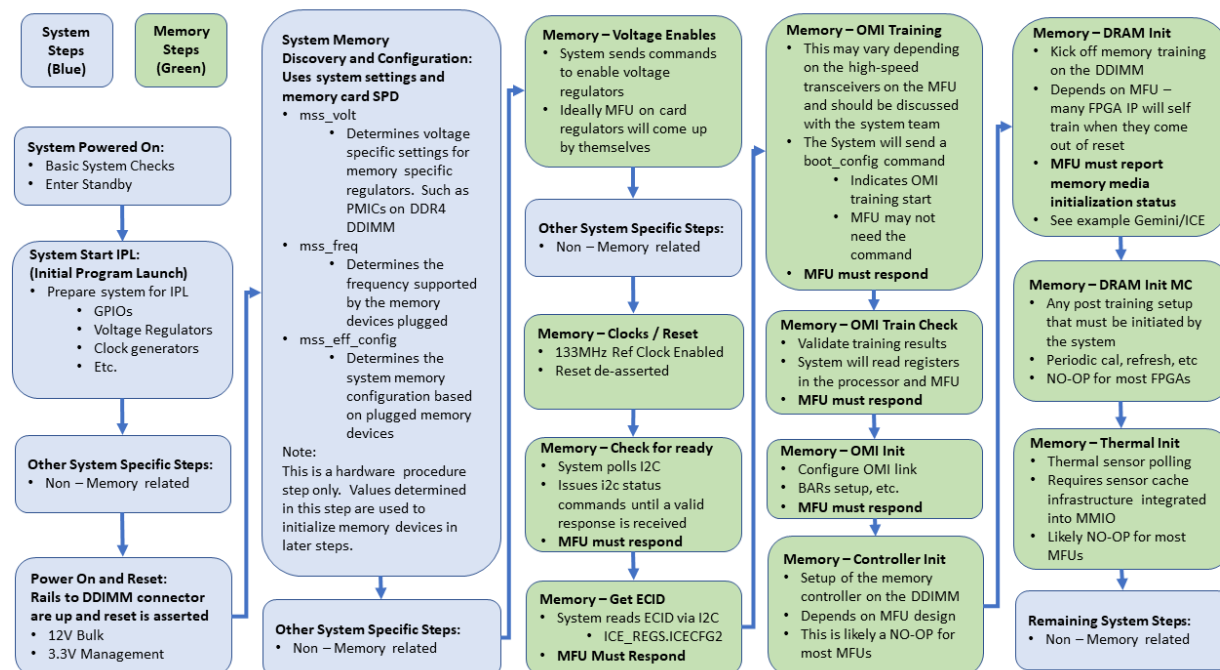
This chapter describes the requirements for interaction with the host firmware in most OMI enabled OpenPOWER systems.

It is outside the scope of this document to discuss every detailed interaction of the host level firmware and software with the deployed MFU as each MFU will vary in its functionality and such variance will likely impact the system software design. However, there are key interactions between the system firmware and the deployed MFU to achieve the basics of allowing the OpenPOWER OMI system to understand what type of memory device is available and how to initialize the device and map its available memory space.

8.1. Firmware Initialization Flow

There are a wide range of potential applications which may deploy an MFU attached to the OPMB MSL. Variants of the MFU may require specific updates to the firmware controlled hardware initialization flow. This section describes the basic high level flow expected to initialize the MFU. This document covers MFUs connected to an OMI channel in an OpenPOWER system, other deployments of OMI and MFU are outside the scope of this document. There are OpenPOWER system variations and it is advised that developers obtain the later initialization flow from their OpenPOWER representative. In addition, some MFUs may require modification to the initialization flow to support specific features, again, it is advised that the developers contact their OpenPOWER representative for support. This document covers only the minimum requirements to initialize an MFU on the OMI channel in an OpenPOWER system.

Figure 8.1. High Level Firmware Initialization Flow



9. Development, Prototyping and Bring-Up Insights

This chapter provides insights and some level of guidance on getting started with development, prototyping and bring-up of your OMI MFU design.

No system required! Don't worry, you don't need a lot of hardware to get going. In fact a RaspberryPi and a reasonable FPGA card can get you going. In fact, that's what we started with when developing the Gemini / ICE design found on Github (see section 1.3). Now it's even easier as most of the tools you will need to ensure compatibility of your design with an OpenPOWER system are available through the OpenPOWER Foundation via the OpenPOWER Toolkit Memtools and through the OpenCAPI consortium.

Depending upon the state of MFU development the OpenPOWER ecosystem can enable various levels of the prototyping process. By no means is deploying an MFU into an OpenPOWER system an all or nothing process. The OpenPOWER Ecosystem enables memory development tools and actual system level hardware procedures to run on stand alone MFU boards. This allows for low cost early development work without the need of a OpenPOWER OMI enabled system. In fact, the example design available on Github (Gemini/Ice in section 1.3 of this document) was developed and tested in a stand alone environment before any OMI systems were available. The example design was also used as a vehicle to test OMI Memory initialization procedures before real system hardware was available. The OpenPOWER Memory Development ecosystem was designed with an understanding that not all components of the design, including both hardware and software would be understood at the beginning and will require an iterative design process. From early prototyping on low cost test benches to final product integration and qualification, the OpenPOWER Ecosystem has it covered.

Cloud development virtual machines	Developer Resources / Software Developer Cloud Resources / <empty>
---	--

Developer Tools	Developer Resources / Developer Tools / <empty>
------------------------	---



Note

Use the **Search** field to focus your search using key words or phrases for specific resources.

A.3. Contact the foundation

To learn more about the OpenPOWER Foundation, please use the following contact points:

- General information -- <info@openpowerfoundation.org>
- Membership -- <membership@openpowerfoundation.org>
- Technical Work Groups and projects -- <tsc-chair@openpowerfoundation.org>
- Events and other activities -- <admin@openpowerfoundation.org>
- Press/Analysts -- <press@openpowerfoundation.org>

More contact information can be found at openpowerfoundation.org/get-involved/contact-us.

VSEC	Vendor-Specific Extended Configuration Capability
VSID	Virtual segment ID.
WIMG bits	Four bits in the page table, also called a page-table entry, which control the processor's accesses to cache and main storage. 'W' stands for write through, 'I' for cache inhibit, 'M' for memory coherence, and 'G' for guarded storage.
word	Four bytes.