



1

2 **Memory-Mapped Buffer Interface (MMBI)**
3 **Specification**

4 **Version: 1.0.2**

5 **Document Identifier: DSP0282**

6 **Date: 2025-10-01**

7 **Version History: <https://www.dmtf.org/dsp/DSP0282>**

8 **Supersedes: 1.0.1**

9 **Document Class: Normative**

10 **Document Status: Published**

11 **Document Language: en-US**

Copyright Notice

Copyright © 2023–2025 DMTF. All rights reserved.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. Members and non-members may reproduce DMTF specifications and documents for uses consistent with this purpose, provided that correct attribution is given. As DMTF specifications may be revised from time to time, the particular version and release date should always be noted.

Implementation of certain elements of this standard or proposed standard may be subject to third-party patent rights, including provisional patent rights (herein “patent rights”). DMTF makes no representations to users of the standard as to the existence of such rights and is not responsible to recognize, disclose, or identify any or all such third-party patent right owners or claimants, nor for any incomplete or inaccurate identification or disclosure of such rights, owners, or claimants. DMTF shall have no liability to any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize, disclose, or identify any such third-party patent rights, or for such party’s reliance on the standard or incorporation thereof in its products, protocols, or testing procedures. DMTF shall have no liability to any party implementing such standards, whether such implementation is foreseeable or not, nor to any patent owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is withdrawn or modified after publication, and shall be indemnified and held harmless by any party implementing the standard from any and all claims of infringement by a patent owner for such implementations.

For information about patents held by third parties which have notified DMTF that, in their opinion, such patents may relate to or impact implementations of DMTF standards, visit <https://www.dmtf.org/about/policies/disclosures>.

PCI-SIG®, PCI Express®, and PCIe® are registered trademarks or service marks of PCI-SIG. All other marks and brands are the property of their respective owners.

This document’s normative language is English. Translation into other languages is permitted.

CONTENTS

39	Foreword	6
40	Introduction	7
41	1 Scope	8
42	2 Normative references	8
43	3 Terms and definitions	8
44	4 Conventions	10
45	4.1 Reserved and unassigned values	10
46	4.2 Byte ordering	10
47	5 Assumptions	10
48	5.1 Underlying Memory Mapping	10
49	5.2 Multiple Instances	10
50	5.3 Resets and Errors	11
51	5.4 Notifications (Interrupts)	11
52	5.5 Packet Sizes, Types, and Packet Flow	11
53	5.6 Security	12
54	6 Basic Architecture Concept	12
55	7 MMBI Data Structures	13
56	7.1 MMBI Capability Descriptor	14
57	7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer	16
58	7.2.1 Variable Packet Size Circular Buffer Descriptor	16
59	7.2.2 Host Read-Write Structure	18
60	7.2.3 Host Read-Only Structure	19
61	8 Runtime Flows	21
62	8.1 MMBI Interface Initialization and Reset	21
63	8.1.1 Initialization of Descriptor Structures after Power Up	21
64	8.1.2 Interface States and Graceful Reset	22
65	8.1.3 Ungraceful Reset Considerations	28
66	8.2 Calculation of Filled Space and Empty Space in Circular Buffer	29
67	8.3 Device Readiness and Communication Pause	29
68	8.4 Packet Transfer	31
69	8.5 Interrupts (Optional)	32
70	9 Multi-Protocol Packet Format	32
71	ANNEX A (informative) Notations	34
72	ANNEX B (informative) Change log	35
73		

74 **Figures**

75	Figure 1 – Multiple MMBI Instances.....	11
76	Figure 2 – MMBI Interface Concept Overview.....	13
77	Figure 3 – MMBI Data Structure Relationships.....	14
78	Figure 4 – MMBI Capability Descriptor Layout	15
79	Figure 5 – MMBI Interface States	25
80	Figure 6 – Sample MMBI Reset by Host.....	26
81	Figure 7 – Sample MMBI Reset Flow by (B)MC	28
82	Figure 8 – Filled and Empty Space in Circular Buffers	29
83	Figure 9 – Sample MMBI Device Pause Sequences.....	30
84		

85 **Tables**

86	Table 1 – MMBI Capability Descriptor Structure (MMBI_Desc).....	15
87	Table 2 – Buffer Type Dependent Descriptor for BUFT=0001b (VPSCB Descriptor)	17
88	Table 3 – MMBI Host Read-Write Structure (Host_RWS).....	19
89	Table 4 – MMBI Host Read-Only Structure (Host_ROS).....	20
90	Table 5 – MMBI Interface States	23
91	Table 6 – Multi-Protocol Packet Format.....	33
92		

93 Foreword

94 The *Memory-Mapped Buffer Interface (MMBI) Specification* (DSP0282) was prepared by the PMCI
95 Working Group of DMTF.

96 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
97 management and interoperability. For more information about DMTF, visit dmf.org.

98 This version supersedes version 1.0.1. For a list of changes, see the change log in ANNEX B.

99 **Acknowledgments**

100 DMTF acknowledges the following individuals for their contributions to this document:

101 Editors:

- 102 • Janusz Jurski – Intel Corporation
- 103 • Richard Marian Thomaiyar – Intel Corporation
- 104 • Jose Marinho – Arm Limited
- 105 • Ramesha He – Dell Inc.

106 Contributors:

- 107 • Rama Bisa – Dell Inc.
- 108 • Patrick Caporale – Lenovo
- 109 • Samer El-Haj-Mahmoud – ARM Inc.
- 110 • Ted Emerson – Hewlett Packard Enterprise
- 111 • John Guan – Inspur
- 112 • Tiffany Kasanicky – Intel Corporation
- 113 • Eliel Louzoun - Intel Corporation
- 114 • Mahesh Natsu – Intel Corporation
- 115 • Chandra Nelogal – Dell Inc.
- 116 • Edward Newman – Hewlett Packard Enterprise
- 117 • Scott Phuong – Cisco
- 118 • Derek Roberts – Xilinx Inc.
- 119 • William Scherer III – Hewlett Packard Enterprise
- 120 • Hemal Shah – Broadcom Inc.
- 121 • Bob Stevens – Dell Inc.

122

Introduction

123 The *Memory-Mapped Buffer Interface (MMBI) Specification* defines the mechanisms facilitating
124 communication between platform components, typically host software and a Management Controller
125 (usually a Baseboard Management Controller). Using the shared memory concept, this document defines
126 the MMBI protocol that allows packet exchanges between communicating devices. The described
127 memory mapping makes it possible for both boot code (such as UEFI firmware), as well as OS-level
128 software (such as an OS kernel or drivers) to establish efficient communication with a (Baseboard)
129 Management Controller at bandwidth and latency limited by the underlying memory mapping
130 mechanisms. MMBI can also be used to enable communication between other types of platform
131 components, not just host software and a Management Controller (MC) or a Baseboard Management
132 Controller (BMC).

1 Scope

This document provides the specifications for the Memory-Mapped Buffer Interface (MMBI). MMBI assumes an underlying memory mapping capability, such as PCIe MMIO/BAR, allowing host software to efficiently access data stored in (B)MC memory. MMBI defines generic packet-based communication mechanism (based on circular buffers), and specific protocols, such as MCTP, should be covered in other documents.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated or versioned references, only the edition cited (including any corrigenda or DMTF update versions) applies. For references without a date or version, the latest published edition of the referenced document (including any corrigenda or DMTF update versions) applies.

DMTF, DSP0236, *Management Component Transport Protocol (MCTP) Base Specification 1.3*,
https://www.dmtf.org/standards/published_documents/DSP0236_1.3.pdf

DMTF, DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.10*,
https://www.dmtf.org/standards/published_documents/DSP0239_1.10.pdf

DMTF, DSP0276, *Secured Messages using SPDm over MCTP Binding Specification 1.1.0*,
https://www.dmtf.org/standards/published_documents/DSP0276_1.1.0.pdf

DMTF, DSP0284, *Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface (MMBI) Transport Binding Specification 1.0*,
https://www.dmtf.org/standards/published_documents/DSP0284_1.0.pdf

IANA, *Internet Assigned Numbers Authority – Private Enterprise Numbers (PEN)*,
<https://www.iana.org/assignments/enterprise-numbers>

PCI-SIG, PCI Express® Base Specification Revision 6.2, February 12, 2024
<https://www.pcisig.com/specifications/>

3 Terms and definitions

In this document, some terms have a specific meaning beyond the normal English meaning. Those terms are defined in this clause.

The terms “shall” (“required”), “shall not”, “should” (“recommended”), “should not” (“not recommended”), “may”, “need not” (“not required”), “can” and “cannot” in this document are to be interpreted as described in [ISO/IEC Directives, Part 2](#), Clause 7. The terms in parentheses are alternatives for the preceding term, for use in exceptional cases when the preceding term cannot be used for linguistic reasons. Note that [ISO/IEC Directives, Part 2](#), Clause 7 specifies additional alternatives. Occurrences of such additional alternatives shall be interpreted in their normal English meaning.

The terms “clause”, “subclause”, “paragraph”, and “annex” in this document are to be interpreted as described in [ISO/IEC Directives, Part 2](#), Clause 6.

The terms “normative” and “informative” in this document are to be interpreted as described in [ISO/IEC Directives, Part 2](#), Clause 3. In this document, clauses, subclauses, or annexes labeled “(informative)” do not contain normative content. Notes and examples are always informative elements.

Refer to [Management Component Transport Protocol \(MCTP\) Base Specification](#) for the terms and definitions that are used across the MCTP specifications.

For the purposes of this document, the following terms and definitions apply.

174	3.1
175	ACK
176	Acknowledge
177	3.2
178	B2H
179	BMC-to-Host
180	3.3
181	BAR
182	Base Address Register
183	3.4
184	(B)MC
185	Baseboard Management Controller – term used interchangeably with Management Controller
186	3.5
187	CCT
188	Control Command Type
189	3.6
190	Destination Device
191	Device receiving the MCTP packet over MMBI
192	3.7
193	H2B
194	Host-to-BMC
195	3.8
196	MMBI
197	Memory-Mapped Buffer Interface
198	3.9
199	MMIO
200	Memory-Mapped Input/Output
201	3.10
202	NACK
203	Not acknowledge
204	3.11
205	ROS
206	Read-Only Structure
207	3.12
208	RWS
209	Read-Write Structure
210	3.13
211	Source Device
212	Device sending the MCTP packet over MMBI

213 **3.14**
214 **SPDM**
215 Security Protocol and Data Model
216 **3.15**
217 **VPSCB**
218 Variable Packet Size Circular Buffer

219 **4 Conventions**

220 The conventions described in the following clauses apply to this specification.

221 **4.1 Reserved and unassigned values**

222 Unless otherwise specified, any reserved, unspecified, or unassigned values in enumerations or other
223 numeric ranges are reserved for future definition by DMTF.

224 Unless otherwise specified, numeric or bit fields that are designated as reserved shall be written as 0
225 (zero) and ignored when read.

226 **4.2 Byte ordering**

227 Unless otherwise specified, byte ordering of multi-byte numeric fields or bit fields is “Big Endian” (that is,
228 the lower byte offset holds the most significant byte, and higher offsets hold less-significant bytes).

229 **5 Assumptions**

230 **5.1 Underlying Memory Mapping**

231 The fundamental assumption in this specification is that there exists an underlying platform mechanism
232 allowing efficient memory sharing between the communicating entities (such as a host and a
233 management controller). PCIe MMIO is an example of such a mechanism. This specification defines the
234 packet transfer protocol on top of this assumed memory mapping layer.

235 Assumptions about the underlying layer are:


- 236 1) Memory mapping shall guarantee an error-free lossless channel.
- 237 2) The size of atomic operations is at least 4 bytes.
- 238 3) The order of operations must be preserved: writes must be visible to the other party in the order
239 they were executed by the sender; reads cannot be prefetched/cached; if interrupts are used,
240 they must also obey the order of operations.

241 **5.2 Multiple Instances**

242 This specification has been designed with the assumption that a single MMBI instance will serve
243 communication between the two communicating entities only (typically host software and management
244 controller firmware components) and so the interface is not shared between multiple communicating
245 entities.

246 Multiple components in the system, e.g., multiple host tenant / software agents communicating to a
247 (B)MC, can be supported using a plurality of MMBI interfaces (each being an independent instance of the

248 interface), located in different memory locations. Such MMBI instances shall operate independently as
 249 shown in Figure 1:



250
 251 **Figure 1 – Multiple MMBI Instances**

252 **5.3 Resets and Errors**

253 MMBI allows lossless communication as well as graceful reset/initialization on request from a
 254 communicating party (in case of a reset of a software entity). However, MMBI does not provide
 255 guaranteed delivery in case of ungraceful resets of the communicating parties. Applications that care
 256 about data loss in such situations shall employ an ACK packet scheme to verify data reception by the
 257 other party and handle the error if ACK is not received.

258 **5.4 Notifications (Interrupts)**

259 MMBI is designed to execute in both interrupt and polling mode.

260 The memory sharing capability may be accompanied by the ability to receive interrupts by the
 261 communicating software entities. MMBI enables discovery and enables use of the optional interrupt
 262 mechanism for efficient data exchange between communicating entities. If interrupts are used, it is
 263 assumed that the interrupt delivery mechanism is reliable.

264 If interrupts are not available, a polling mode can be used. Platform designers can choose polling or
 265 interrupt mode, based on their needs.

266 **5.5 Packet Sizes, Types, and Packet Flow**

267 MMBI allows variable packet sizes, with the maximum size dependent on the underlying physical layer's
 268 memory mapping capabilities. MMBI provides a discovery method allowing the communicating parties to
 269 define and discover the circular buffer sizes, which limit the maximum packet sizes that can be

270 transmitted (fragmentation/reassembly is not supported by this version of MMBI protocol). The upper
271 layers must adhere to the discovered limits and, if necessary, handle fragmentation/reassembly
272 accordingly.

273 MMBI allows multiple packets (datagrams) to be in-flight. That is, the sender can place more than one
274 packet in the memory buffer even before they are consumed by the receiver. This enables asynchronous
275 operation of the communicating entities. Regardless of the number of packets in-flight, they are
276 guaranteed to arrive to the receiver in the FIFO order (note: upper layer can elect to process in same
277 order or in different order, which will not be guaranteed by the MMBI layer). Note that if multiple instances
278 of MMBI are in the system, they operate independently and no packet ordering guarantees exist between
279 them.

280 MMBI enables and defines discovery mechanisms to support the exchange of a variety of packet protocol
281 types, such as MCTP. Binding of these protocols to MMBI is defined in separate documents, such as
282 [*Management Component Transport Protocol \(MCTP\) Memory-Mapped Buffer Interface \(MMBI\) Transport*](#)
283 [*Binding Specification*](#).

284 5.6 Security

285 MMBI does not provide any security guarantees. Any authentication, integrity protection, and/or
286 encryption is to be implemented by the other layers of the protocol stack. For example, for secure
287 implementation of communication between the host and (B)MC using MMBI, [*Secured Messages using*](#)
288 [*SPDM over MCTP Binding Specification*](#) can be used. Another alternative can be host-based memory
289 protection mechanisms.

290 6 Basic Architecture Concept

291 The host and the (B)MC use circular buffers to exchange data. One buffer is used to send data from the
292 host to the (B)MC and is referred to as H2B (Host-to-BMC). The other buffer is used for communication in
293 the opposite direction and is referred to as B2H (BMC-to-Host). The buffers are used to store packet data,
294 and they are accompanied by a descriptor structure. The descriptor is a data structure in the shared
295 memory used to store important capabilities and control information. These data structures are shown in
296 Figure 2 and are defined in detail in section 7.




Figure 2 – MMBI Interface Concept Overview

7 MMBI Data Structures

Each instance of the MMBI interface is divided into sections as defined below:

- “BMC-to-Host” (B2H) region with substructure as follows:
 - *MMBI Capability Descriptor (MMBI_Desc Structure)* — see section 7.1 for details
 - *Host_ROS* (Host Read-Only Structure) — see section 7.2.3 for details
 - (B)MC-to-Host Circular buffer (B2H Circular buffer) — see section 8 for details
- “Host-to-BMC” (H2B) region with substructure as follows:
 - *Host_RWS* (Host Read-Write Structure) — see section 7.2.2 for details
 - Host-to-(B)MC circular buffer (H2B Circular buffer) — see section 8 for details

The format of the H2B and B2H circular buffers is a sequence of packets, and this format is referred to as Variable Packet Size Circular Buffer (VPSCB). For VPSCB, the relationships between these data structures and their main pointers are as presented in Figure 3.




Figure 3 – MMBI Data Structure Relationships

Details of these data structures are presented in the following subsections. Note that the data structures maintain 4-byte alignment for fields that need to be updated atomically. Packets in the circular buffers are also aligned to 4-byte boundaries.

7.1 MMBI Capability Descriptor

MMBI Capability Descriptor is used to define the MMBI interface details. (B)MC updates this data structure during initialization. Other than that, the (B)MC and host are not allowed to update it. The host only reads this descriptor to understand the format of the MMBI data structures in memory and shall never write to this data structure. The layout of the structure is presented in Figure 4 and described in Table 1. See also section 8.1.




Figure 4 – MMBI Capability Descriptor Layout

Table 1 – MMBI Capability Descriptor Structure (MMBI_Desc)

Byte(s)	Description
0:5	MMBI Signature “#MMBI\$” in ASCII. When this signature is not present, the host SW should assume the absence of MMBI.
6	[7:4] Reserved [3:0] MMBI version 0001b – Implementations of MMBI described in this document shall indicate version 1 of MMBI.
7	[7:1] Reserved [0] OS Use Indicates if this MMBI interface is intended for OS use: 0b – OS should not use this MMBI interface as it is managed by other host software components (UEFI BIOS, ACPI ASL code, etc.). 1b – This MMBI interface is intended for OS use.
8:11	[31:29] – Reserved [28:0] B2H Buffer Base Address (B2H_BA) B2H (BMC-to-Host) buffer base address expressed in 8-byte units as offset relative to the beginning of the descriptor

Byte(s)	Description
12:15	[31:29] – Reserved
	[28:0] H2B Buffer Base Address (H2B_BA) H2B (Host-to-BMC) buffer base address expressed in 8-byte units as offset relative to the beginning of the descriptor
16:19	B2H Buffer Length (B2H_L) The size of the B2H buffer (can represent up to 4GB)
20:23	H2B Buffer Length (H2B_L) The size of the B2H buffer (can represent up to 4GB)
24	[7:4] Reserved
	[3:0] Buffer Type (BUFT) Indicates the type of data structures in H2B and B2H buffers. The following values are defined: 0001b – MMBI Variable Packet Size Circular Buffers (VPSCB) v1 (see section 7.2) Other values are reserved.
25:31	Reserved
32:52	Buffer Type Dependent Descriptor The definition of this field is dependent on the BUFT field value: If BUFT=0001b (VPSCB), Table 2 in section 7.2 defines the format of these bytes and the packet format in circular buffers is defined in section 9
56:63	Reserved

326

327 7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer

328 This section describes data structures used when the communication between (B)MC and host SW
 329 happens according to the VPSCB Buffer Type (BUFT=0001b).

330 7.2.1 Variable Packet Size Circular Buffer Descriptor

331 Variable Packet Size Circular Buffer Descriptor is part of the *MMBI_Desc* structure. Its access rules are
 332 the same as *MMBI_Desc*:

- 333 • The (B)MC updates this data structure during MMBI interface initialization.
- 334 • Neither the (B)MC nor the host are allowed to update it at any other time.

335

Table 2 – Buffer Type Dependent Descriptor for BUFT=0001b (VPSCB Descriptor)

Byte(s)	Description
0:3	[31:29] – Reserved
	[28:0] Host Read-Only Structure Pointer (H_ROS_P) Points to the <i>Host_ROS</i> structure. The base address is expressed in 8-byte units as the offset relative to beginning of the descriptor
4:7	[31:29] – Reserved
	[28:0] Host Read-Write Structure Pointer (H_RWS_P) Points to the <i>Host_RWS</i> structure. The base address is expressed in 8-byte units as the offset relative to beginning of the descriptor
8	[7:3] – Reserved
	[2:0] Host Interrupt Type (H_Int_T) Defines how the (B)MC interrupts the host. This is an informative field from the host's perspective with the intention to keep the (B)MC and host in sync. 0 – no interrupt / polling 1 – PCIe interrupt (bus specific) 2 – physical pin (GPIO) 3 – eSPI Virtual Wire Other values are reserved
9	Host Interrupt Location (H_Int_L) If H_Int_T = 0: reserved If H_Int_T = 1: for PCIe, indicates the PCIe interrupt message number If H_Int_T = 2: pin number If H_Int_T = 3: eSPI Virtual Wire Index number Reserved otherwise
10:12	Reserved
13	Host Interrupt Value (H_Int_V) If H_Int_T = 3: eSPI Virtual Wire data value Reserved otherwise

Byte(s)	Description
14	[7:3] – Reserved [2:0] (B)MC Interrupt Type (BMC_Int_T) Defines how the (B)MC wants to be interrupted: 0 – no interrupt triggering by the host 1 – relative memory space address (offset defined in the BMC_Int_L field) 2 – Inband interrupt (bus specific—such as PCIe MSI or virtual legacy wire) Other values – reserved
15:18	(B)MC Interrupt Location (BMC_Int_L) If BMC_Int_T = 1, memory address—offset relative to the beginning of the <i>MMBI Capability Descriptor</i> base address Otherwise reserved
19:22	Reserved
23	(B)MC Interrupt Value (BMC_Int_V) If BMC_Int_T = 1, this field indicates the value to be written at the given address to trigger an interrupt. Otherwise reserved

7.2.2 Host Read-Write Structure

The host's RW Structure Pointer in the above structure points to the *Host_RWS* structure, which is shown in Table 1. This structure is accessed as follows:

- It is initialized by the (B)MC to the default values.
- The host updates this structure during normal communication—it is read-writeable for the host.
- The (B)MC is not allowed to write to this structure during normal communication—it should treat this structure as read-only (any kind of hardware-based enforcement of the read-only behavior is out of scope of this specification).

344

Table 3 – MMBI Host Read-Write Structure (Host_RWS)

Byte(s)	Description
0:3	[31:2] H2B Write Pointer (H2B_WP) Bits [31:2] of the offset where the host can write the next data in the H2B circular buffer, counted from the beginning of the H2B buffer represented in 4-byte alignment. Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment). The (B)MC uses this pointer to determine how many bytes of valid data are present in the Circular Buffer (by comparing it with the H2B_RP offset). The host shall advance the pointer once data is written to the Circular Buffer and shall update this pointer to mark the next available offset. Note: The host shall not overwrite the data not read by the (B)MC, as indicated by the H2B_RP.
	[1] Host Interface Up (H_UP) 1 indicates that the host side of the interface is up and running, which means that the data structures can be used by the (B)MC.
	[0] Host Reset Request (H_RST) Setting this flag to 1 will initiate a reset sequence to get the circular buffers into a known good state (see section 8.1 for more information).
4:7	[31:2] B2H Read Pointer (B2H_RP) Bits [31:2] of the offset where the host reads data from the B2H circular buffer, counted from the beginning of the B2H buffer represented in 4-byte alignment. Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment). The (B)MC uses this pointer to determine how much of data is read by the host. Comparing this with the B2H Write Pointer (B2H_WP) will provide how much space is left to write the data. The host shall only advance the pointer once the data available in B2H is read by the host.
	[1] Reserved
	[0] Host Ready (H_RDY) 0 indicates that the host is performing some tasks that keep it busy, and so it may be unresponsive. However, the (B)MC can use the data structures and, for example, put data into the buffers as long as H_UP = 1. 1 indicates that the host is ready to exchange data (see section 8.1 for more information).

345 7.2.3 Host Read-Only Structure

346 Host RO Structure Pointer points to *Host_ROS* structure. The host is only allowed to read this structure
 347 (never write). Any kind of hardware-based enforcement of the read-only behavior is out-of-scope of this
 348 specification. This structure is initialized by the (B)MC to the default values and later updated by (B)MC
 349 during normal communication—it is read-writeable for the (B)MC.

350

Table 4 – MMBI Host Read-Only Structure (Host_ROS)

Byte(s)	Description
0:3	[31:2] B2H Write Pointer (B2H_WP) Bits [31:2] of the offset where the (B)MC can write the next data in the B2H circular buffer, counted from the beginning of the B2H buffer. Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment). The host uses this pointer to determine how many bytes of valid data are present in the Circular Buffer (by comparing it with B2H_RP offset) The (B)MC shall advance the pointer once data is written to the Buffer to mark the next available offset. Note: (B)MC shall not overwrite the data not read by host, as indicated by the B2H_RP.
	[1] (B)MC Interface Up (B_UP) 1 indicates that the (B)MC side of the interface is up and running which means that the data structures are initialized and can be used
	[0] (B)MC Reset Request (B_RST) Setting this flag to 1 will initiate a reset sequence to get the circular buffers into a known good state (see section 8.1 for more information).
4:7	[31:2] H2B Read Pointer (H2B_RP) Bits [31:2] of the offset where the host reads data from the H2B circular buffer, counted from the beginning of the H2B buffer. Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment). The host uses this pointer to determine how much of data is read by the (B)MC. Comparing this with the H2B write pointer will provide how much space is left to write. (B)MC shall only advance the pointer once the data available in H2B is read by the (B)MC.
	[1] Reserved
	[0] (B)MC Ready (B_RDY) 0 indicates that the (B)MC is performing some tasks that keep it busy and so it may be unresponsive – host however can use the data structures and, for example, put data into the buffers as long as B_UP = 1 1 indicates that the (B)MC is ready to exchange data (see section 8.1 for more information).

351 MMBI uses two circular buffers: H2B and B2H. Each buffer is a memory range defined in the descriptor
 352 with the following access:

- 353 • H2B (Host-to-BMC buffer) is RW for the host and RO for the (B)MC.
- 354 • B2H (BMC-to-Host buffer) is RO for the host and RW for the (B)MC.

355 The Read Pointer and Write Pointer are used to indicate the read and write location in the buffer. For
 356 each read or write the pointer shall be advanced. It means pointer increment with a rollover at the buffer
 357 size.

358 These pointers, along with the Buffer Length fields (B2H_L or H2B_L), are used to calculate the number
 359 of filled bytes to read or the number of empty bytes available for write.

360 The circular buffers will be used to send packets of arbitrary size. A packet may require multiple memory
 361 reads and/or write transfers.

362 8 Runtime Flows

363 8.1 MMBI Interface Initialization and Reset

364 This section describes the steps to allow the (B)MC to complete the initialization of the data structures
 365 and indicating when both sides of communication are ready to exchange data.

366 The goal of the reset, on the other hand, is to reinitialize the data structures when at least one side wants
 367 a clean start, which may be due to unexpected device events, malfunction, error, etc. It may also be used
 368 to reinitialize the data structures after, for example, a (B)MC firmware update in which the data structure
 369 needs some new values (e.g., when the circular buffer size changes after the firmware update). A
 370 graceful reset follows the state diagram presented in Figure 5, and it guarantees that MMBI protocol layer
 371 does not drop any packets (note that other protocol layers may still be unable to guarantee delivery).

372 The reset sequence is also automatically initiated when hardware errors lead to all-ones or all-zeros
 373 memory reads, as is typical with some media. This is thanks to the fact that when all the flags are zeros or
 374 are all ones, it indicates an initialization or transition to initialization states. Such unexpected resets do not
 375 follow the handshake protocol, and so are ungraceful and may lead to packet losses.

376 These flags are used to indicate the (B)MC's status as related to initialization and reset:

- 377 • (B)MC Interface Up (B_UP)
- 378 • (B)MC Reset Request (B_RST)

379 Similar flags are used to indicate the host's status:

- 380 • Host Interface Up (H_UP)
- 381 • Host Reset Request (H_RST)

382 All these flags are used in combination to achieve the proper handshake mechanism between the host
 383 and the (B)MC during initialization or reset.

384 8.1.1 Initialization of Descriptor Structures after Power Up

385 The (B)MC must initialize the expected content of the MMBI data structures (see section 7) during power
 386 up and make the shared memory available to the host. Initialization is expected to complete before the
 387 host software accesses these structures so that the host can find the *MMBI Capability Descriptor*
 388 (*MMBI_Desc*) using the MMBI signature bytes. MMBI structures and buffers must always remain
 389 available in the shared memory when the host is using the MMBI interface.

390 If the MMBI is made available via a memory-mapped range of a PCIe function, then the *MMBI_Desc* is at
 391 offset 0 of a PCIe function's BARs, and the function's PCIe Base Class, Sub-Class, and Programming
 392 Interface shall be {0xC, 0xC, 0x0}. There can be at most one *MMBI_Desc* per BAR.

393 During the initial accesses after the host's power up or reset, the host's software is expected to verify if
394 the content of the MMBI version and MMBI signature are as expected. If the above requirements are met,
395 the host is expected to check the interface state.

396 If the host's software does not find the proper *MMBI Capability Descriptor (MMBI_Desc)* content at the
397 expected location, the host should consider the MMBI as not present or, optionally, it may implement a
398 wait option with a timeout. Such a timeout mechanism is system-dependent and is out of scope of this
399 specification.

400 If the MMBI signature and MMBI version fields match, but the size and location of the buffers cannot be
401 fulfilled by the host, it shall indicate the initialization mismatch error by transitioning to the *Initialization*
402 *Mismatch* state as described below. With this indication, the (B)MC may consider the interface as
403 inoperable or attempt to reinitialize the *MMBI_Desc* structure with, for example, a smaller buffer size.
404 Before updating the data structure content, the (B)MC shall first clear the B_UP flag and then clear the
405 H_RST flag to return back to the *Initialization in Progress* state. Such attempts to repair the situation are
406 system-dependent and are out of scope of this specification.

407 8.1.2 Interface States and Graceful Reset

408 When _RST and _UP are both set on one side of communication, it means the entity is requesting a reset
409 sequence. When B_RST = H_RST = B_UP = H_UP = 1, it means that both entities are ready to perform
410 the reset sequence (in fact, the host is just waiting for the (B)MC to do all the initialization).

411 All the states are summarized in Table 5. The "Host Write Access" and "(B)MC Write Access" columns
412 define write-access restrictions to the data structures by host and (B)MC, respectively. There are no read
413 restrictions for the (B)MC and host. Note that the host is expected to re-read the data structure contents
414 after initialization is completed.


Table 5 – MMBI Interface States

B_UP	B_RST	H_UP	H_RST	State Description	Host Write Access	(B)MC Write Access
0	0	0	0	Initialization in Progress The (B)MC is initializing the data structures. The host can only monitor the data structures, waiting for B_UP = 1 and B_RST = 0 flags.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
1	0	0	0	Initialization Completed The (B)MC has completed initialization of the data structures and is ready to exchange data—waiting for the host to be ready. The host should re-read the <i>MMBI_Desc</i> structure and any dependent structures. During this state, the (B)MC is allowed to deposit packets into the circular buffer.	Host allowed to write to MMBI structures as per section 7	(B)MC allowed to write to MMBI structures as per section 7
1	0	1	0	Normal Runtime Both the (B)MC and host use the data structures and the circular buffers for data exchanges.	Host allowed to write to MMBI structures as per section 7	(B)MC allowed to write to MMBI structures as per section 7
1	1	1	0	Reset Request by (B)MC The (B)MC is requesting reset—waiting for the host to notice the request. When the host notices the request, it should consume the data from the B2H (if any) and shall set H_RST flag as an ACK and wait for the initialization to complete (B_UP = 1 and B_RST = 0 status).	Host allowed to write to MMBI structures as per section 7	(B)MC allowed to write to MMBI structures as per section 7
1	0	1	1	Reset Request by Host The host is requesting reset—waiting for the (B)MC to notice the request and reinitialize the interface. When the host sets the H_RST flag, it shall not perform any further updates in the MMBI data structures but shall only wait for the initialization to be completed by (B)MC (B_UP = 1 and B_RST = 0 status). When the (B)MC notices the request, it should consume the data from the B2H (if any) and shall set B_RST flag as an ACK.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures

B_UP	B_RST	H_UP	H_RST	State Description	Host Write Access	(B)MC Write Access
1	1	1	1	Reset ACKed The host and (B)MC are ready to perform graceful interface reset. This is a transient state when the host is waiting for the (B)MC to complete the initialization. The host is not allowed to write to MMBI data structures. The (B)MC is expected to clear the B_UP and B_RST flags (in this order) and reinitialize all the data structures.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
0	1	1	1	Transitioning to Initialization Transient state after the “Reset ACKed” state. The host is not allowed to write to MMBI data structures.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
0	1	1	0	Temporary Transition States These states may be observed during initialization when the (B)MC updates the data structures (reinitialization of all the data structures is not an atomic operation). They are unexpected during normal operation and if they happen it means that MMBI structures have been corrupted. The (B)MC may initialize the interface or stop using MMBI and report a fatal error.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
0	1	0	1			
0	1	0	0			
0	0	0	1			
1	0	0	1	Initialization Mismatch The host causes transition into this state from <i>Initialization Completed</i> when it is unable to use the interface due to unsupported content in the <i>MMBI Capability Descriptor</i> structure.	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
1	1	0	1	Unexpected States		
1	1	0	0	These states shall never happen:		
0	0	1	0	If the (B)MC reads this state, it indicates that the host does not follow MMBI protocol or some other corruption happened—the (B)MC should initialize the interface or it may stop using MMBI and report a fatal error, depending on system policy.		
0	0	1	1	If the host reads this state, it may wait for the reinitialization to complete or stop using MMBI and report a fatal error, depending on system policy.		

416 The expected state transitions are presented in Figure 5:

417



418

419

Figure 5 – MMBI Interface States

420 The host shall check the MMBI Interface state before writing any new data to the H2B buffer (as
 421 described in Table 5, the host is only allowed to transfer new data in the Normal Runtime state, i.e.,
 422 B_UP=1 & B_RST=0 & H_UP=1 & H_RST=0). Similarly, the (B)MC shall check the status before writing

any new data to the B2H buffer. These status flags are conveniently located in the B2H_WP or H2B_WP bytes which the host or (B)MC, respectively, read anyway during any use of the circular buffers.

8.1.2.1 Host Initiating Graceful Reset Sequence

Assuming *Normal Runtime* state, the host shall use the following sequence to request MMBI interface reset:

- 1) The host sets H_RST = 1 to initiate the reset flow. If (B)MC interrupts are enabled, the host notifies the (B)MC.
 - a. In response, the (B)MC is expected to set B_RST = 1, which indicates the transition to the *Reset ACKed* state. If host interrupts are enabled, the host is expected to be notified about the update (or else it uses polling). At this point, the (B)MC reinitializes all the data structures.
- 2) The host waits for B_UP = 1 and B_RST = 0 (and H_UP = H_RST = 0), which indicates the transition to the *Initialization Completed* state. Host interrupts are not used at this stage until H_UP is set by host software.
- 3) The host transitions to the *Normal Runtime* state by setting H_UP = 1. The host is also expected to set the B_RDY flag, indicating that it can receive and handle new packets—see section 8.3. If (B)MC interrupts are enabled, the host notifies the (B)MC after the flags are updated.

Figure 6 presents a sample flow:




Figure 6 – Sample MMBI Reset by Host

8.1.2.2 (B)MC Initiating Graceful Reset Sequence

Assuming *Normal Runtime* state, the (B)MC shall use the following sequence to request MMBI interface reset:

- 1) The (B)MC sets B_RST = 1 to initiate the reset flow. If host interrupts are enabled, the (B)MC notifies the host.
- 2) The (B)MC waits for H_UP = 1 and H_RST = 1, which indicates the transition to the *Reset ACKed* state. If (B)MC interrupts are enabled, the (B)MC is expected to be notified about the update (or else (B)MC uses polling).
- 3) The (B)MC clears the B_UP flag (B_RST still set). Host interrupts are no longer enabled.
- 4) The (B)MC clears the H_UP and H_RST flags (this may cause transient states to be observed by the host).
- 5) The (B)MC clears the B_RST flag.
- 6) The (B)MC reinitializes all the data structures.
- 7) The (B)MC sets B_UP = 1. Host interrupts are not used at this stage until H_UP is set by host software.
- 8) The (B)MC waits for the host to set H_UP = 1. If (B)MC interrupts are enabled, the (B)MC is expected to be notified about the update.

Note that the (B)MC is also expected to set the B_RDY flag, typically in step 7, indicating that it can receive and handle new packets—see section 8.3.

Figure 7 presents a sample flow.




Figure 7 – Sample MMBI Reset Flow by (B)MC

8.1.3 Ungraceful Reset Considerations

If an ungraceful reset/crash happens, MMBI does not guarantee delivery. However, provisions are present in the MMBI design to handle the following scenarios:

1. In the case of a (B)MC FW-only reset (HW continues to work, memory content, including buffers stay intact in shared memory and accesses are still handled by HW): the host will still see the MMBI in the normal state and write to MMBI Circular buffers to deposit or read data as long as there is any space available in the buffers. In this situation, host may timeout waiting for a response but this is handled by higher layers above MMBI.
2. (B)MC HW reset (buffers are wiped and MMIO mechanisms are broken): the host will see errors on reads/writes and must handle them as per host-specific mechanisms. Additionally, MMBI encoding of status in B_UP, B_RST, H_UP, & H_RST is such that all-zeros or all-ones are recognized as transient states (see Table 5). So, even if there would be no other mechanisms in the system, the host would still recognize this as an error and would have to wait for reinitialization by the (B)MC (the host is not allowed to write to the buffers in the transient state, i.e., until the data structures are reinitialized by (B)MC FW).
3. Unexpected host reset (SW or HW reset is the same outcome): the host's unexpected reset will leave the data structures intact in (B)MC memory, so the (B)MC can still read the data from the buffers. Assuming the (B)MC understands the host's status via other mechanisms, the (B)MC can take informed decisions about how to respond to such situations.

In all the above cases, MMBI data structures can be reinitialized after the reset to allow a clean restart.

8.2 Calculation of Filled Space and Empty Space in Circular Buffer

The procedure for calculating the number of filled bytes in a circular buffer is analogous for both the H2B and B2H buffers: the difference between the write pointer and read pointer indicates the amount of valid data, accounting for the rollover at the end of the buffer. The write pointer cannot advance beyond the read pointer, accounting for the rollover at the end of the buffer.

The following steps allow calculation of the number of filled slots in a circular buffer:

1. The write and read pointers must start with zero after initialization. Since read pointer = write pointer, there is no valid data/packets in the buffer on initialization.
2. Once data is written to the buffer, the source (the host or (B)MC) will advance the write buffer pointer.
3. Read pointer is advanced once data is read/consumed by the receiver (the host or (B)MC).
4. Rollover: when the pointers reach the maximum offset within the buffer during writing/reading, data must be written/read starting back at zero offset, and the pointers roll over accordingly.




Figure 8 – Filled and Empty Space in Circular Buffers

8.3 Device Readiness and Communication Pause

In addition to the reinitialization or reset states, the MMBI interface also uses the H_RDY and B_RDY flags to indicate the device's readiness to consume incoming packets and handle them. When the host or (B)MC are ready to receive and handle packets, they set the B_RDY or H_RDY flags, respectively. If a B_RDY or H_RDY flag is clear but the B_UP and H_UP flags are set, it means that the MMBI interface is up but the target device is not ready to consume and handle new packets. When the interface is up, it means that the data structures are ready to accept new packets so the sender can:

- wait for the receiver to become ready before writing new packets to the buffer—this is important if the sender expects an action to be taken by the receiver, such as providing a response

- deposit new packets to the buffer in order for the receiver to consume them later—this capability may be used if the sender does not expect a response from the receiver; for example, when the sender needs to deposit some logs in the shared memory

An example flow when (B)MC firmware / host software undergoes a reset and indicates its non-readiness during a reboot is presented in Figure 9. Note that in this example it is assumed that the MMBI data structures are still intact in shared memory during the reset.




Figure 9 – Sample MMBI Device Pause Sequences

8.4 Packet Transfer

This flow describes the host-to-(B)MC flow that shall be followed to send a packet. An analogous flow shall be followed to send packets in the opposite direction (swap (B)MC and host in the description and use B2H buffer instead):

- 1) Host software reads the read and write pointers (H2B_RP and H2B_WP) to determine the number of empty spaces available in its circular buffer.
 - a. If there is not enough empty space available in the host's circular buffer, the host waits until there is room in the host's circular buffer. This is done either by polling or waiting for an interrupt.
 - b. Host software shall also verify that B_UP = 1 & B_RST = 0 before any packet transfers. If this is not the case, it shall follow the reset process as defined in section 8.1. Note that the host may decide to delay packet transfer depending on B_RDY state and its policy.
- 2) Once there is enough empty space available in the circular buffer, the host writes the packet into the host's circular buffer. To accomplish this, the host sequentially writes data at the write pointer location but not exceeding the length of the buffer (H2B_L). When it reaches the maximum address of the buffer, it shall continue writing the packet from the buffer base address (H2B_BA). This process shall never overflow the buffer by advancing beyond the H2B_RP.
- 3) Once the packet write is complete, the host updates the write pointer value in H2B_WP.
- 4) In Interrupt-enabled mode, the (B)MC firmware is interrupted:
 - a. If BMC_Int_T = 1, the host uses the (B)MC Interrupt Info (BMC_Int_L) and (B)MC Interrupt Value (BMC_Int_V) to interrupt the (B)MC FW.
 - b. Even if BMC_Int_T = 0, the (B)MC HW may also monitor H2B_WP and generate an interrupt automatically.
 - c. Alternatively, a platform-specific method can be used to trigger the interrupt to (B)MC.
- 5) In polling mode, the (B)MC FW can continuously read the write pointer to see when it changes. In interrupt mode, it is woken up by the (B)MC HW.
- 6) The (B)MC firmware reads the read/write pointers and determines the number of filled spaces in the circular buffer available for reading.
 - a. If the circular buffer is empty, the host has not sent a packet. This interrupt is for another reason, or it indicates that the host has completed reading the packet(s) last transmitted by the (B)MC.
- 7) The (B)MC FW reads the buffer data written by the host.
- 8) The (B)MC FW updates the read pointer in B2H_RWS. This indicates to the host how much data has been read by the (B)MC, and the host can use the portion of the buffer that has been read already.
- 9) If host notifications are enabled, (B)MC FW shall generate an interrupt to the host.
- 10) When the host software gets interrupted or due to polling of H2B_RP, it can determine that the (B)MC has consumed the data. The host can also poll instead of relying on interrupts.

8.5 Interrupts (Optional)

Interrupts, if enabled by the discovery/control mechanisms of MMBI, shall be triggered for the following reasons (both for host software and (B)MC firmware):

- A packet has just been written to the circular buffer.
- A packet has just been read from the circular buffer.
- The host or (B)MC has initiated an interface reset sequence.
- The host or (B)MC has completed its portion of the interface reset sequence and normal operation can begin.

An interrupt handler shall:

- check the status flags in the *MMBI Capability Descriptor (MMBI_Desc)*—if a reset is initiated, the flow defined in section 8.1 shall be followed
- check if there is a packet in the circular buffer—this can be calculated as per section 8.2—and, if there is data present in the buffer, the interrupt handler should initiate the packet receive flow, as defined in section 8.4.

If there are multiple instances of the MMBI interface sharing the same interrupt, the interrupt handler shall check all the instances for the reasons listed above. The order of such a check and interrupt affinity are implementation-specific and out of scope of this specification.

9 Multi-Protocol Packet Format

If BUFT=0001b (VPSCB) and Packet Protocol Type = 0001b (Multi-protocol Type), the multi-protocol MMBI packets will have the following defined header fields, as shown in Table 6. There is a 4-byte alignment expectation, meaning that padding must be added if necessary for the packet length to be a multiple of 4 bytes.

578

Table 6 – Multi-Protocol Packet Format

Byte(s)	Description						
0:2	[23:2] Packet Length (PKT_LEN) The size of the packet, calculated as PKT_LEN+1 multiplied by 4 bytes (can represent up to 16MB packet). Values 0x3FFFFFF and zero are reserved.						
	[1:0] Packet padding (PKT_PAD) Number of padding bytes						
3	[7:4] – Reserved						
	[3:0] – Packet type (PKT_TYPE) Defines the format of the remaining bytes: 0100b – MCTP over MMBI (see Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface (MMBI) Transport Binding Specification) 0101b – Vendor defined content as defined below Other values are reserved.						
4:N-1	Protocol type specific fields This field depends on the PKT_TYPE value: If PKT_TYPE = MCTP = 0100b: format follows MCTP over MMBI (see Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface (MMBI) Transport Binding Specification) If PKT_TYPE = Vendor defined = 0101b: the following vendor-defined format shall be used: <table border="1" data-bbox="402 1100 1349 1289"> <thead> <tr> <th>Byte(s)</th><th>Description</th></tr> </thead> <tbody> <tr> <td>4:7</td><td>Vendor IANA Enterprise Number encoded in little-endian format; for more information about IANA Enterprise Numbers, please see Internet Assigned Numbers Authority – Private Enterprise Numbers</td></tr> <tr> <td>8:N-1</td><td>Content defined by the vendor</td></tr> </tbody> </table>	Byte(s)	Description	4:7	Vendor IANA Enterprise Number encoded in little-endian format; for more information about IANA Enterprise Numbers, please see Internet Assigned Numbers Authority – Private Enterprise Numbers	8:N-1	Content defined by the vendor
Byte(s)	Description						
4:7	Vendor IANA Enterprise Number encoded in little-endian format; for more information about IANA Enterprise Numbers, please see Internet Assigned Numbers Authority – Private Enterprise Numbers						
8:N-1	Content defined by the vendor						
(N:M)	Padding (PAD) – optional Padding bytes as defined in PKT_PAD field. Note: padding is added to ensure packets are 4-byte aligned						

579

ANNEX A (informative)

Notations

Examples of notations used in this document are as follows:

- 2:N In field descriptions, this will typically be used to represent a range of byte offsets starting from byte two and continuing to and including byte N. The lowest offset is on the left; the highest is on the right.
- (6) Parentheses around a single number can be used in packet field descriptions to indicate a byte field that may be present or absent.
- (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range may be present or absent. The lowest offset is on the left; the highest is on the right.
- [DSP0236](#) Underlined blue text is typically used to indicate a reference to a document or specification called out in [Normative references](#) or to items hyperlinked within the document.
- [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets are given as zero-based values (that is, the least significant bit offset = 0).
- [7:5] A range of bit offsets. The most significant bit is on the left, the least significant bit is on the right.
- 1b A number consisting of 0s and 1s followed by a lowercase “b” indicates that the number is in binary format.
- 0x12A A leading “0x” indicates that the number is in hexadecimal format.

ANNEX B
(informative)

Change log

Version	Date	Description
1.0.0	2023-07-14	Initial release
1.0.1	2024-08-30	Document title change ("Memory-Mapped BMC Interface" to "Memory-Mapped Buffer Interface") to better reflect potential broader uses of MMBI beyond just BMC
1.0.2	2025-10-01	Define the discovery process of an MMBI description within a PCIe endpoint. Fixed: https://github.com/DMTF/PMCI-WG/issues/1646