

1



2

3

Document Identifier: DSP0282

4

Date: 2024-10-09

5

Version: 1.0.1

6

# Memory-Mapped Buffer Interface (MMBI)

7

## Specification

8

**Supersedes: 1.0.0**

9

**Document Class: Normative**

10

**Document Status: Published**

11

**Document Language: en-US**

12 Copyright Notice

13 Copyright © 2023–2024 DMTF. All rights reserved.

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

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29 withdrawn or modified after publication, and shall be indemnified and held harmless by any party  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

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

35 All other marks and brands are the property of their respective owners.

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

37

# CONTENTS

38 Foreword ..... 6

39 Introduction..... 7

40 1 Scope ..... 8

41 2 Normative references ..... 8

42 3 Terms and definitions ..... 8

43 4 Conventions ..... 10

44 4.1 Reserved and unassigned values ..... 10

45 4.2 Byte ordering..... 10

46 5 Assumptions ..... 10

47 5.1 Underlying Memory Mapping ..... 10

48 5.2 Multiple Instances ..... 10

49 5.3 Resets and Errors ..... 11

50 5.4 Notifications (Interrupts)..... 11

51 5.5 Packet Sizes, Types, and Packet Flow..... 11

52 5.6 Security ..... 12

53 6 Basic Architecture Concept ..... 12

54 7 MMBI Data Structures ..... 13

55 7.1 MMBI Capability Descriptor ..... 14

56 7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer..... 16

57 7.2.1 Variable Packet Size Circular Buffer Descriptor ..... 16

58 7.2.2 Host Read-Write Structure..... 18

59 7.2.3 Host Read-Only Structure..... 19

60 8 Runtime Flows..... 21

61 8.1 MMBI Interface Initialization and Reset ..... 21

62 8.1.1 Initialization of Descriptor Structures after Power Up ..... 21

63 8.1.2 Interface States and Graceful Reset..... 22

64 8.1.3 Ungraceful Reset Considerations ..... 28

65 8.2 Calculation of Filled Space and Empty Space in Circular Buffer..... 29

66 8.3 Device Readiness and Communication Pause ..... 29

67 8.4 Packet Transfer..... 31

68 8.5 Interrupts (Optional) ..... 32

69 9 Multi-Protocol Packet Format..... 32

70 ANNEX A (informative) Notations ..... 34

71 ANNEX B (informative) Change log..... 35

72

73 **Figures**

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

84 **Tables**

85 Table 1 – MMBI Capability Descriptor Structure (MMBI\_Desc)..... 15

86 Table 2 – Buffer Type Dependent Descriptor for BUFT=0001b (VPSCB Descriptor) ..... 17

87 Table 3 – MMBI Host Read-Write Structure (Host\_RWS) ..... 19

88 Table 4 – MMBI Host Read-Only Structure (Host\_ROS)..... 20

89 Table 5 – MMBI Interface States ..... 23

90 Table 6 – Multi-Protocol Packet Format..... 33

91

92

## Foreword

93 The *Memory-Mapped Buffer Interface (MMBI) Specification* (DSP0282) was prepared by DMTF's PMCI  
94 Working Group.

95 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems  
96 management and interoperability. For information about DMTF, visit [dmtf.org](http://dmtf.org).

97 This version supersedes version 1.0.0. For a list of changes, see the change log in ANNEX B.

## 98 Acknowledgments

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

100 Editors:

- 101 • Janusz Jurski – Intel Corporation
- 102 • Richard Marian Thomaiyar – Intel Corporation

103 DMTF Contributors:

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

119

## Introduction

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

## 130 1 Scope

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

## 136 2 Normative references

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

141 DMTF, DSP0236, *Management Component Transport Protocol (MCTP) Base Specification 1.3*,  
142 [https://www.dmtf.org/standards/published\\_documents/DSP0236\\_1.3.pdf](https://www.dmtf.org/standards/published_documents/DSP0236_1.3.pdf)

143 DMTF, DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.10*,  
144 [https://www.dmtf.org/standards/published\\_documents/DSP0239\\_1.10.pdf](https://www.dmtf.org/standards/published_documents/DSP0239_1.10.pdf)

145 DMTF, DSP0276, *Secured Messages using SPDm over MCTP Binding Specification 1.1.0*,  
146 [https://www.dmtf.org/standards/published\\_documents/DSP0276\\_1.1.0.pdf](https://www.dmtf.org/standards/published_documents/DSP0276_1.1.0.pdf)

147 DMTF, DSP0284, *Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface  
148 (MMBI) Transport Binding Specification 1.0*,  
149 [https://www.dmtf.org/standards/published\\_documents/DSP0284\\_1.0.pdf](https://www.dmtf.org/standards/published_documents/DSP0284_1.0.pdf)

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

## 152 3 Terms and definitions

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

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

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

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

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

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



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

208 **3.14**  
209 **SPDM**  
210 Security Protocol and Data Model

211 **3.15**  
212 **VPSCB**  
213 Variable Packet Size Circular Buffer

## 214 **4 Conventions**

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

### 216 **4.1 Reserved and unassigned values**

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

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

### 221 **4.2 Byte ordering**

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

## 224 **5 Assumptions**

### 225 **5.1 Underlying Memory Mapping**

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

230 Assumptions about the underlying layer are:

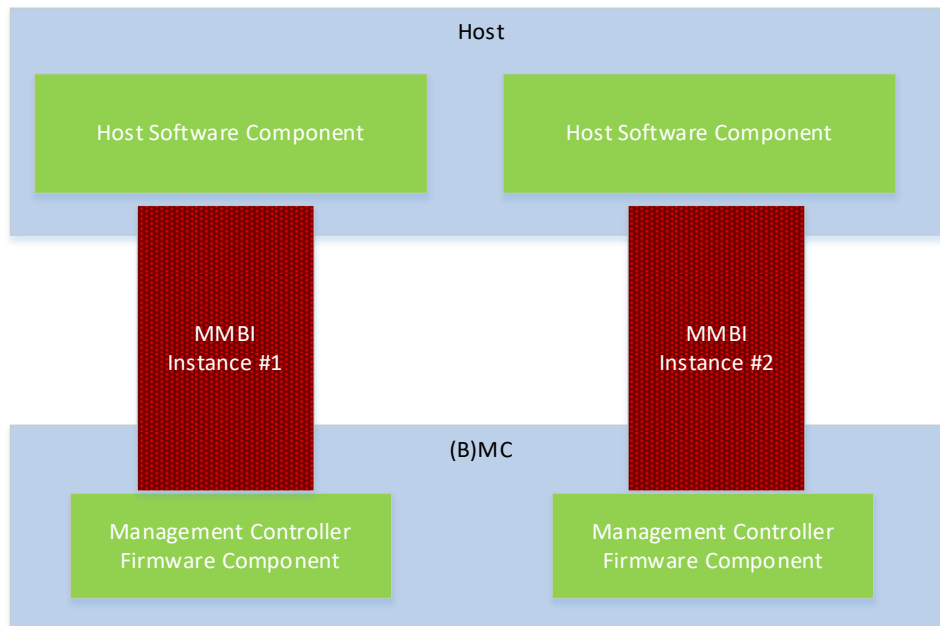
- 231 1) Memory mapping shall guarantee an error-free lossless channel.
- 232 2) The size of atomic operations is at least 4 bytes.
- 233 3) The order of operations must be preserved: writes must be visible to the other party in the order  
234 they were executed by the sender; reads cannot be prefetched/cached;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

### 236 **5.2 Multiple Instances**

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

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

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



245

246

**Figure 1 – Multiple MMBI Instances**

247 **5.3 Resets and Errors**

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

253 **5.4 Notifications (Interrupts)**

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

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

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

261 **5.5 Packet Sizes, Types, and Packet Flow**

262 MMBI allows variable packet sizes, with the maximum size dependent on the underlying physical layer's  
 263 memory mapping capabilities. MMBI provides a discovery method allowing the communicating parties to  
 264 define and discover the circular buffer sizes, which limit the maximum packet sizes that can be  
 265 transmitted (fragmentation/reassembly is not supported by this version of MMBI protocol). The upper

266 layers must adhere to the discovered limits and, if necessary, handle fragmentation/reassembly  
267 accordingly.

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


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

## 279 5.6 Security

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

## 285 6 Basic Architecture Concept

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



292

293

**Figure 2 – MMBI Interface Concept Overview**


## 294 7 MMBI Data Structures

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

- 296 • “BMC-to-Host” (B2H) region with substructure as follows:
  - 297 ○ *MMBI Capability Descriptor (MMBI\_Desc Structure)* — see section 7.1 for details
  - 298 ○ *Host\_ROS (Host Read-Only Structure)* — see section 7.2.3 for details
  - 299 ○ *BMC-to-Host Circular buffer (B2H Circular buffer)* — see section 8 for details
- 300 • “Host-to-BMC” (H2B) region with substructure as follows:
  - 301 ○ *Host\_RWS (Host Read-Write Structure)* — see section 7.2.2 for details
  - 302 ○ *Host-to-BMC circular buffer (H2B Circular buffer)* — see section 8 for details

303

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



307


308

**Figure 3 – MMBI Data Structure Relationships**

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

312 **7.1 MMBI Capability Descriptor**

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



318

319

Figure 4 – MMBI Capability Descriptor Layout

320

Table 1 – MMBI Capability Descriptor Structure (MMBI\_Desc)

Byte(s)	Description
0:5	<b>MMBI Signature</b> “#MMBI\$” in ASCII. When this signature is not present, the host SW should assume the absence of MMBI.
6	[7:4] Reserved
	<b>[3:0] MMBI version</b> 0001b – Implementations of MMBI described in this document shall indicate version 1 of MMBI.
7	[7:1] Reserved
	<b>[0] OS Use</b> 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
	<b>[28:0] B2H Buffer Base Address (B2H_BA)</b> 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
	<b>[28:0] H2B Buffer Base Address (H2B_BA)</b> H2B (Host-to-BMC) buffer base address expressed in 8-byte units as offset relative to the beginning of the descriptor
16:19	<b>B2H Buffer Length (B2H_L)</b> The size of the B2H buffer (can represent up to 4GB)
20:23	<b>H2B Buffer Length (H2B_L)</b> The size of the B2H buffer (can represent up to 4GB)
24	[7:4] Reserved
	<b>[3:0] Buffer Type (BUFT)</b> 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	<b>Buffer Type Dependent Descriptor</b> 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

321

## 322 7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer

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

### 325 7.2.1 Variable Packet Size Circular Buffer Descriptor

326 Variable Packet Size Circular Buffer Descriptor is part of the *MMBI\_Desc* structure. Its access rules are  
327 the same as *MMBI\_Desc*:

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



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

Byte(s)	Description
0:3	[31:29] – Reserved
	<p><b>[28:0] Host Read-Only Structure Pointer (H_ROS_P)</b></p> <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</p>
4:7	[31:29] – Reserved
	<p><b>[28:0] Host Read-Write Structure Pointer (H_RWS_P)</b></p> <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</p>
8	[7:3] – Reserved
	<p><b>[2:0] Host Interrupt Type (H_Int_T)</b></p> <p>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.</p> <p>0 – no interrupt / polling                      1 – PCIe interrupt (bus specific)                      2 – physical pin (GPIO)                      3 – eSPI Virtual Wire                      Other values are reserved</p>
9	<p><b>Host Interrupt Location (H_Int_L)</b></p> <p>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</p>
10:12	Reserved
13	<p><b>Host Interrupt Value (H_Int_V)</b></p> <p>If H_Int_T = 3: eSPI Virtual Wire data value                      Reserved otherwise</p>

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

### 331 7.2.2 Host Read-Write Structure

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

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

339

**Table 3 – MMBI Host Read-Write Structure (Host\_RWS)**

Byte(s)	Description
0:3	<p><b>[31:2] H2B Write Pointer (H2B_WP)</b></p> <p>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.</p> <p>Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment).</p> <p>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).</p> <p>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.</p> <p>Note: The host shall not overwrite the data not read by the (B)MC, as indicated by the H2B_RP.</p>
	<p><b>[1] Host Interface Up (H_UP)</b></p> <p>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.</p>
	<p><b>[0] Host Reset Request (H_RST)</b></p> <p>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).</p>
4:7	<p><b>[31:2] B2H Read Pointer (B2H_RP)</b></p> <p>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.</p> <p>Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment).</p> <p>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.</p> <p>The host shall only advance the pointer once the data available in B2H is read by the host.</p>
	<p><b>[1] Reserved</b></p>
	<p><b>[0] Host Ready (H_RDY)</b></p> <p>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.</p> <p>1 indicates that the host is ready to exchange data (see section 8.1 for more information).</p>

340 **7.2.3 Host Read-Only Structure**

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

345

Table 4 – MMBI Host Read-Only Structure (Host\_ROS)

Byte(s)	Description
0:3	<p><b>[31:2] B2H Write Pointer (B2H_WP)</b></p> <p>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.</p> <p>Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment).</p> <p>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)</p> <p>The (B)MC shall advance the pointer once data is written to the Buffer to mark the next available offset.</p> <p>Note: (B)MC shall not overwrite the data not read by host, as indicated by the B2H_RP.</p>
	<p><b>[1] (B)MC Interface Up (B_UP)</b></p> <p>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</p>
	<p><b>[0] (B)MC Reset Request (B_RST)</b></p> <p>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).</p>
4:7	<p><b>[31:2] H2B Read Pointer (H2B_RP)</b></p> <p>Bits [31:2] of the offset where the host reads data from the H2B circular buffer, counted from the beginning of the H2B buffer.</p> <p>Bits [1:0] of the offset are assumed to always be zero (for 4-byte alignment).</p> <p>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.</p> <p>(B)MC shall only advance the pointer once the data available in H2B is read by the (B)MC.</p>
	<p><b>[1] Reserved</b></p>
	<p><b>[0] (B)MC Ready (B_RDY)</b></p> <p>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</p> <p>1 indicates that the (B)MC is ready to exchange data (see section 8.1 for more information).</p>

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

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

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

353 These pointers, along with the Buffer Length fields (B2H\_L or H2B\_L), are used to calculate the number  
354 of filled bytes to read or the number of empty bytes available for write.

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

## 357 **8 Runtime Flows**

### 358 **8.1 MMBI Interface Initialization and Reset**

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

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

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

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

- 372 • (B)MC Interface Up (B\_UP)
- 373 • (B)MC Reset Request (B\_RST)

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

- 375 • Host Interface Up (H\_UP)
- 376 • Host Reset Request (H\_RST)

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

#### 379 **8.1.1 Initialization of Descriptor Structures after Power Up**

380 The (B)MC must initialize the expected content of the MMBI data structures (see section 7) during power  
381 up and make the shared memory available to the host. Initialization is expected to complete before the  
382 host software accesses these structures so that the host can find the *MMBI Capability Descriptor*  
383 (*MMBI\_Desc*) using the MMBI signature bytes. MMBI structures and buffers must always remain  
384 available in the shared memory when the host is using the MMBI interface.

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

388 If the host's software does not find the proper *MMBI Capability Descriptor (MMBI\_Desc)* content at the  
389 expected location, the host should consider the MMBI as not present or, optionally, it may implement a

390 wait option with a timeout. Such a timeout mechanism is system-dependent and is out of scope of this  
391 specification.

392 If the MMBI signature and MMBI version fields match, but the size and location of the buffers cannot be  
393 fulfilled by the host, it shall indicate the initialization mismatch error by transitioning to the *Initialization*  
394 *Mismatch* state as described below. With this indication, the (B)MC may consider the interface as  
395 inoperable or attempt to reinitialize the *MMBI\_Desc* structure with, for example, a smaller buffer size.  
396 Before updating the data structure content, the (B)MC shall first clear the B\_UP flag and then clear the  
397 H\_RST flag to return back to the *Initialization in Progress* state. Such attempts to repair the situation are  
398 system-dependent and are out of scope of this specification.

### 399 **8.1.2 Interface States and Graceful Reset**

400 When \_RST and \_UP are both set on one side of communication, it means the entity is requesting a reset  
401 sequence. When B\_RST = H\_RST = B\_UP = H\_UP = 1, it means that both entities are ready to perform  
402 the reset sequence (in fact, the host is just waiting for the (B)MC to do all the initialization).

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

Table 5 – MMBI Interface States

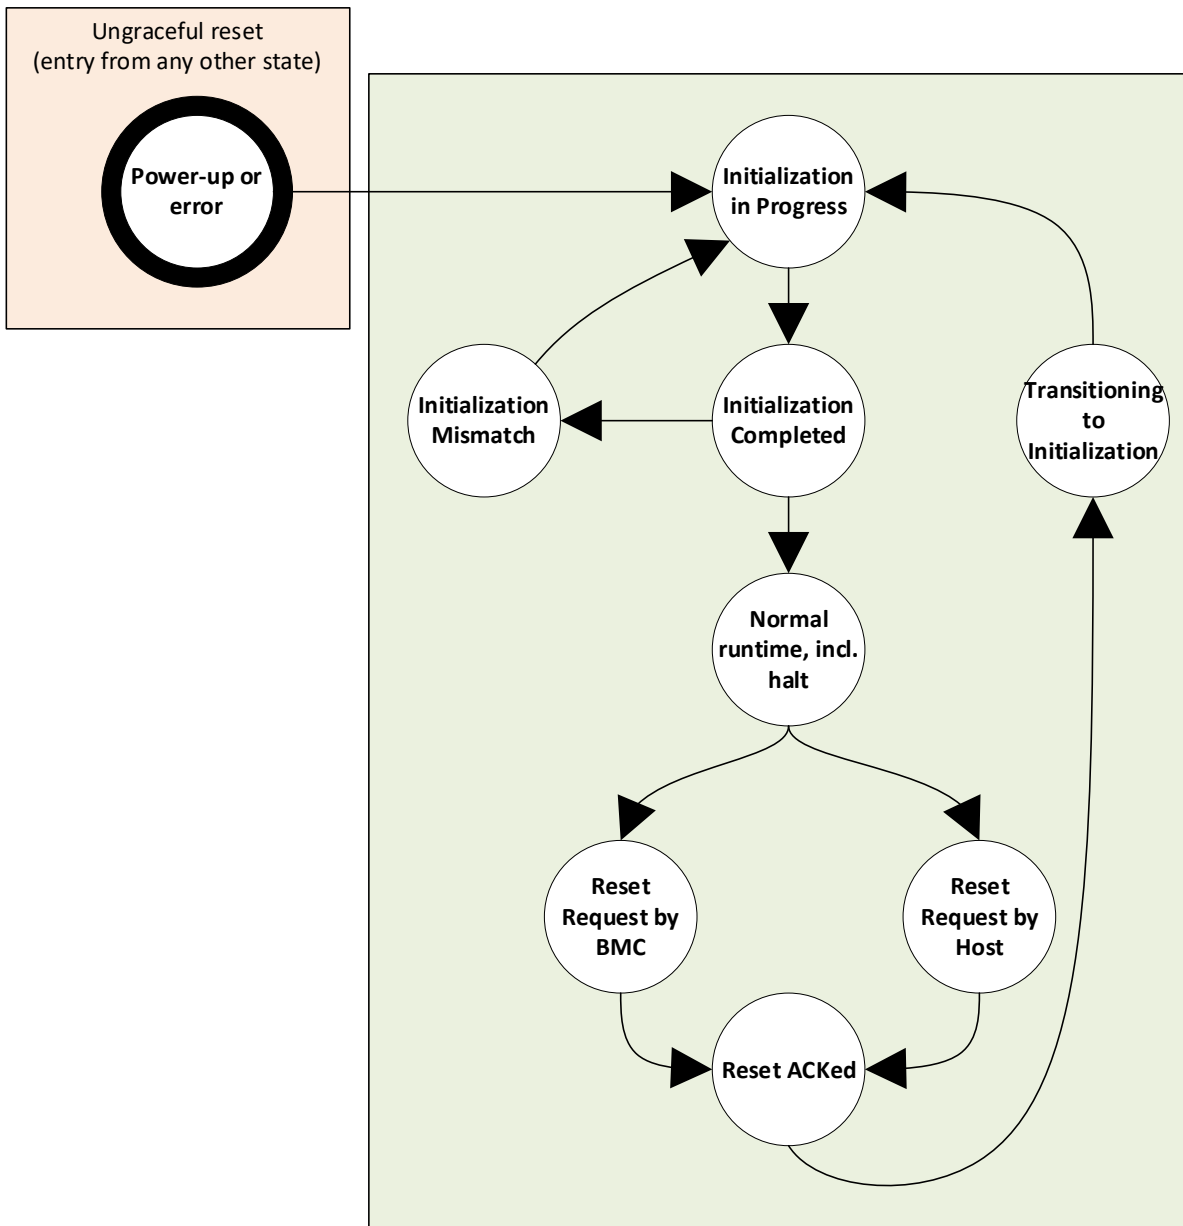
B_UP	B_RST	H_UP	H_RST	State Description	Host Write Access	(B)MC Write Access
0	0	0	0	<p><b>Initialization in Progress</b></p> <p>The (B)MC is initializing the data structures.</p> <p>The host can only monitor the data structures, waiting for B_UP = 1 and B_RST = 0 flags.</p>	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
1	0	0	0	<p><b>Initialization Completed</b></p> <p>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.</p> <p>During this state, the (B)MC is allowed to deposit packets into the circular buffer.</p>	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	<p><b>Normal Runtime</b></p> <p>Both the (B)MC and host use the data structures and the circular buffers for data exchanges.</p>	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	<p><b>Reset Request by (B)MC</b></p> <p>The (B)MC is requesting reset—waiting for the host to notice the request.</p> <p>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).</p>	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	<p><b>Reset Request by Host</b></p> <p>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).</p> <p>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.</p>	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	<b>Reset ACKed</b> 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	<b>Transitioning to Initialization</b> 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	<b>Temporary Transition States</b> 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	<b>Initialization Mismatch</b> 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	<b>Unexpected States</b> These states shall never happen: <ul style="list-style-type: none"> <li>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.</li> <li>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.</li> </ul>		
1	1	0	0			
0	0	1	0			
0	0	1	1			



408 The expected state transitions are presented in Figure 5:

409



410

411

**Figure 5 – MMBI Interface States**

412 The host shall check the MMBI Interface state before writing any new data to the H2B buffer (as  
 413 described in Table 5, the host is only allowed to transfer new data in the Normal Runtime state, i.e.,  
 414 B\_UP=1 & B\_RST=0 & H\_UP=1 & H\_RST=0). Similarly, the (B)MC shall check the status before writing


415 any new data to the B2H buffer. These status flags are conveniently located in the B2H\_WP or H2B\_WP  
 416 bytes which the host or (B)MC, respectively, read anyway during any use of the circular buffers.

417 **8.1.2.1 Host Initiating Graceful Reset Sequence**

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

- 420 1) The host sets H\_RST = 1 to initiate the reset flow. If (B)MC interrupts are enabled, the host  
 421 notifies the (B)MC.
  - 422 a. In response, the (B)MC is expected to set B\_RST = 1, which indicates the transition to  
 423 the *Reset ACKed* state. If host interrupts are enabled, the host is expected to be notified  
 424 about the update (or else it uses polling). At this point, the (B)MC reinitializes all the data  
 425 structures.
- 426 2) The host waits for B\_UP = 1 and B\_RST = 0 (and H\_UP = H\_RST = 0), which indicates the  
 427 transition to the *Initialization Completed* state. Host interrupts are not used at this stage until  
 428 H\_UP is set by host software.
- 429 3) The host transitions to the *Normal Runtime* state by setting H\_UP = 1. The host is also expected  
 430 to set the B\_RDY flag, indicating that it can receive and handle new packets—see section 8.3. If  
 431 (B)MC interrupts are enabled, the host notifies the (B)MC after the flags are updated.

432 Figure 6 presents a sample flow:



433

434

**Figure 6 – Sample MMBI Reset by Host**

435 **8.1.2.2 (B)MC Initiating Graceful Reset Sequence**

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

438 1) The (B)MC sets B\_RST = 1 to initiate the reset flow. If host interrupts are enabled, the (B)MC  
439 notifies the host.

440 2) The (B)MC waits for H\_UP = 1 and H\_RST = 1, which indicates the transition to the *Reset ACKed*  
441 state. If (B)MC interrupts are enabled, the (B)MC is expected to be notified about the update (or  
442 else (B)MC uses polling).

443 3) The (B)MC clears the B\_UP flag (B\_RST still set). Host interrupts are no longer enabled.

444 4) The (B)MC clears the H\_UP and H\_RST flags (this may cause transient states to be observed by  
445 the host).

446 5) The (B)MC clears the B\_RST flag.

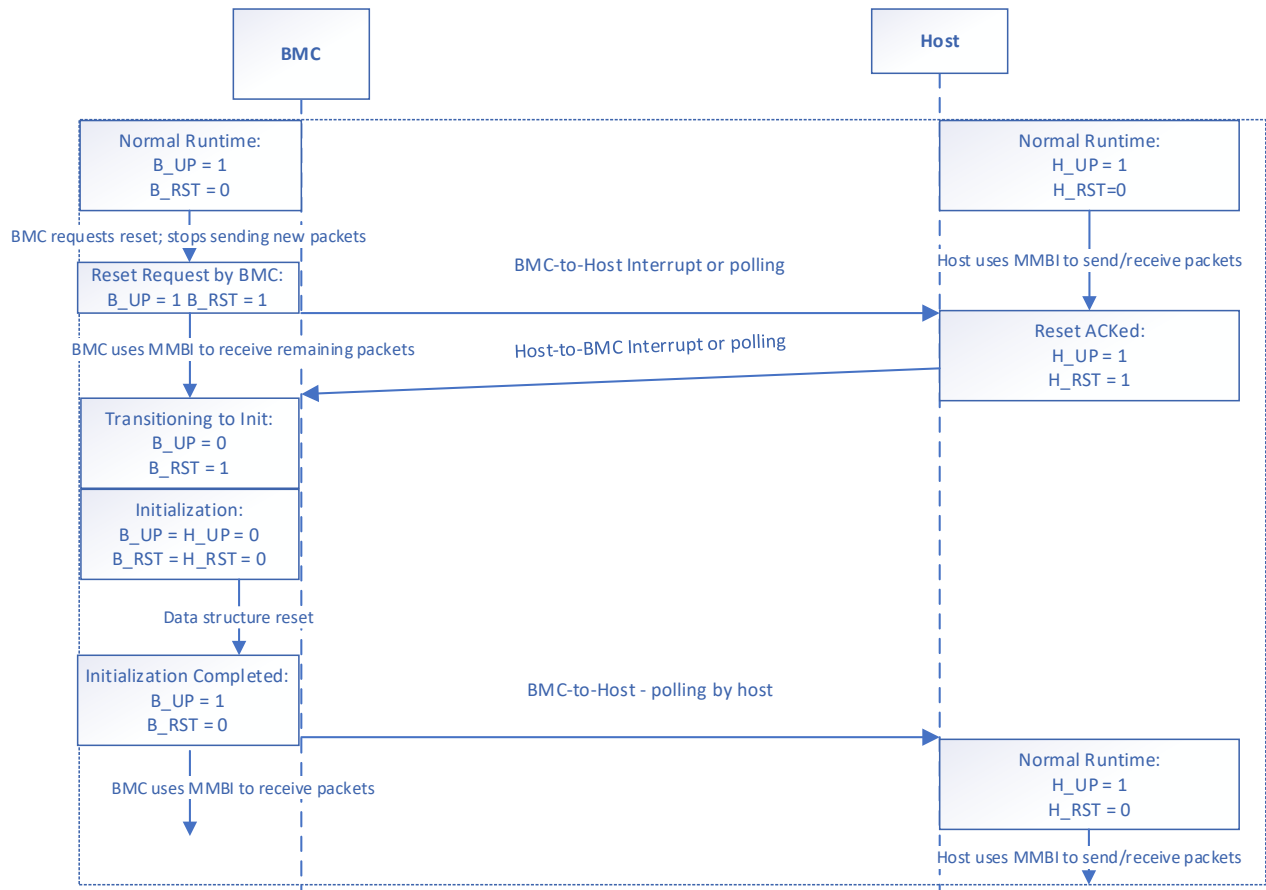
447 6) The (B)MC reinitializes all the data structures.

448 7) The (B)MC sets B\_UP = 1. Host interrupts are not used at this stage until H\_UP is set by host  
449 software.

450 8) The (B)MC waits for the host to set H\_UP = 1. If (B)MC interrupts are enabled, the (B)MC is  
451 expected to be notified about the update.

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

454 Figure 7 presents a sample flow.



455

456

Figure 7 – Sample MMBI Reset by (B)MC

### 457 8.1.3 Ungraceful Reset Considerations

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

- 460
- 461 1. In the case of a (B)MC FW-only reset (HW continues to work, memory content, including  
462 buffers stay intact in shared memory and accesses are still handled by HW): the host will still  
463 see the MMBI in the normal state and write to MMBI Circular buffers to deposit or read data  
464 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.
  - 465 2. (B)MC HW reset (buffers are wiped and MMIO mechanisms are broken): the host will see  
466 errors on reads/writes and must handle them as per host-specific mechanisms. Additionally,  
467 MMBI encoding of status in B\_UP, B\_RST, H\_UP, & H\_RST is such that all-zeros or all-ones  
468 are recognized as transient states (see Table 5). So, even if there would be no other  
469 mechanisms in the system, the host would still recognize this as an error and would have to  
470 wait for reinitialization by the (B)MC (the host is not allowed to write to the buffers in the  
471 transient state, i.e., until the data structures are reinitialized by (B)MC FW).
  - 472 3. Unexpected host reset (SW or HW reset is the same outcome): the host's unexpected reset  
473 will leave the data structures intact in (B)MC memory, so the (B)MC can still read the data  
474 from the buffers. Assuming the (B)MC understands the host's status via other mechanisms,  
475 the (B)MC can take informed decisions about how to respond to such situations.


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

477 **8.2 Calculation of Filled Space and Empty Space in Circular Buffer**

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

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

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



490

491 **Figure 8 – Filled and Empty Space in Circular Buffers**


492 **8.3 Device Readiness and Communication Pause**

493 In addition to the reinitialization or reset states, the MMBI interface also uses the H\_RDY and B\_RDY  
 494 flags to indicate the device’s readiness to consume incoming packets and handle them. When the host or  
 495 (B)MC are ready to receive and handle packets, they set the B\_RDY or H\_RDY flags, respectively. If a  
 496 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  
 497 up but the target device is not ready to consume and handle new packets. When the interface is up, it  
 498 means that the data structures are ready to accept new packets so the sender can:

- 499 • wait for the receiver to become ready before writing new packets to the buffer—this is  
 500 important if the sender expects an action to be taken by the receiver, such as providing a  
 501 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.



508

509

Figure 9 – Sample MMBI Device Pause Sequences

## 510 8.4 Packet Transfer

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

- 514 1) Host software reads the read and write pointers (H2B\_RP and H2B\_WP) to determine the  
515 number of empty spaces available in its circular buffer.
  - 516 a. If there is not enough empty space available in the host's circular buffer, the host waits  
517 until there is room in the host's circular buffer. This is done either by polling or waiting for  
518 an interrupt.
  - 519 b. Host software shall also verify that B\_UP = 1 & B\_RST = 0 before any packet transfers. If  
520 this is not the case, it shall follow the reset process as defined in section 8.1. Note that  
521 the host may decide to delay packet transfer depending on B\_\_RDY state and its policy.
- 522 2) Once there is enough empty space available in the circular buffer, the host writes the packet into  
523 the host's circular buffer. To accomplish this, the host sequentially writes data at the write pointer  
524 location but not exceeding the length of the buffer (H2B\_L). When it reaches the maximum  
525 address of the buffer, it shall continue writing the packet from the buffer base address (H2B\_BA).  
526 This process shall never overflow the buffer by advancing beyond the H2B\_RP.
- 527 3) Once the packet write is complete, the host updates the write pointer value in H2B\_WP.
- 528 4) In Interrupt-enabled mode, the (B)MC firmware is interrupted:
  - 529 a. If BMC\_Int\_T = 1, the host uses the (B)MC Interrupt Info (BMC\_Int\_L) and (B)MC  
530 Interrupt Value (BMC\_Int\_V) to interrupt the (B)MC FW.
  - 531 b. Even if BMC\_Int\_T = 0, the (B)MC HW may also monitor H2B\_WP and generate an  
532 interrupt automatically.
  - 533 c. Alternatively, a platform-specific method can be used to trigger the interrupt to (B)MC.
- 534 5) In polling mode, the (B)MC FW can continuously read the write pointer to see when it changes. In  
535 interrupt mode, it is woken up by the (B)MC HW.
- 536 6) The (B)MC firmware reads the read/write pointers and determines the number of filled spaces in  
537 the circular buffer available for reading.
  - 538 a. If the circular buffer is empty, the host has not sent a packet. This interrupt is for another  
539 reason, or it indicates that the host has completed reading the packet(s) last transmitted  
540 by the (B)MC.
- 541 7) The (B)MC FW reads the buffer data written by the host.
- 542 8) The (B)MC FW updates the read pointer in B2H\_RWS. This indicates to the host how much data  
543 has been read by the (B)MC, and the host can use the portion of the buffer that has been read  
544 already.
- 545 9) If host notifications are enabled, (B)MC FW shall generate an interrupt to the host.
- 546 10) When the host software gets interrupted or due to polling of H2B\_RP, it can determine that the  
547 (B)MC has consumed the data. The host can also poll instead of relying on interrupts.

## 548 8.5 Interrupts (Optional)

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

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

556 An interrupt handler shall:

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

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

## 565 9 Multi-Protocol Packet Format

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



570

**Table 6 – Multi-Protocol Packet Format**

Byte(s)	Description						
0:2	<p><b>[23:2] Packet Length (PKT_LEN)</b></p> <p>The size of the packet, calculated as PKT_LEN+1 multiplied by 4 bytes (can represent up to 16MB packet).</p> <p>Values 0x3FFFFFF and zero are reserved.</p>						
	<p><b>[1:0] Packet padding (PKT_PAD)</b></p> <p>Number of padding bytes</p>						
3	<p>[7:4] – Reserved</p>						
	<p><b>[3:0] – Packet type (PKT_TYPE)</b></p> <p>Defines the format of the remaining bytes:</p> <p>0100b – MCTP over MMBI (see <a href="#">Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface (MMBI) Transport Binding Specification</a>)</p> <p>0101b – Vendor defined content as defined below</p> <p>Other values are reserved.</p>						
4:N-1	<p><b>Protocol type specific fields</b></p> <p>This field depends on the PKT_TYPE value:</p> <p>If PKT_TYPE = MCTP = 0100b: format follows MCTP over MMBI (see <a href="#">Management Component Transport Protocol (MCTP) Memory-Mapped Buffer Interface (MMBI) Transport Binding Specification</a>)</p> <p>If PKT_TYPE = Vendor defined = 0101b: the following vendor-defined format shall be used:</p> <table border="1" style="margin-left: 40px;"> <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 <a href="#">Internet Assigned Numbers Authority – Private Enterprise Numbers</a></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 <a href="#">Internet Assigned Numbers Authority – Private Enterprise Numbers</a>	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 <a href="#">Internet Assigned Numbers Authority – Private Enterprise Numbers</a>						
8:N-1	Content defined by the vendor						
(N:M)	<p><b>Padding (PAD) – optional</b></p> <p>Padding bytes as defined in PKT_PAD field.</p> <p>Note: padding is added to ensure packets are 4-byte aligned</p>						

571

## ANNEX A (informative)

### Notations

572  
573  
574  
575  
576

577 Examples of notations used in this document are as follows:

- 578 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets  
579 starting from byte two and continuing to and including byte N. The lowest offset is on  
580 the left; the highest is on the right.
- 581 • (6) Parentheses around a single number can be used in packet field descriptions to  
582 indicate a byte field that may be present or absent.
- 583 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range  
584 may be present or absent. The lowest offset is on the left; the highest is on the right.
- 585 • [PCle](#) Underlined blue text is typically used to indicate a reference to a document or  
586 specification called out in clause 2, "Normative References" or to items hyperlinked  
587 within the document.
- 588 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets  
589 are given as zero-based values (that is, the least significant bit offset = 0).
- 590 • [7:5] A range of bit offsets. The most significant bit is on the left, the least significant bit is  
591 on the right.
- 592 • 1b A number consisting of 0s and 1s followed by a lowercase "b" indicates that the  
593 number is in binary format.
- 594 • 0x12A A leading "0x" indicates that the number is in hexadecimal format.

595  
596  
597  
598  
599

## ANNEX B (informative)

### Change log

Version	Date	Description
1.0.0	2023-07-14	Initial release
1.0.1	2024-10-09	Document title change ("Memory-Mapped BMC Interface" to "Memory-Mapped Buffer Interface") to better reflect potential broader uses of MMBI beyond just BMC

600