

1



2

3

Document Identifier: DSP0282

4

Date: 2023-08-25

5

Version: 1.0.0

6

Memory-Mapped BMC Interface (MMBI)

7

Specification

8

Supersedes: None

9

Document Class: Normative

10

Document Status: Published

11

Document Language: en-US

12 Copyright Notice

13 Copyright © 2023 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, provided that correct attribution is given. As DMTF specifications may be revised from time to
17 time, the particular version and release date should always be noted.

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28 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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

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

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

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

CONTENTS

| | | |
|----|---|----|
| 37 | Foreword | 5 |
| 38 | Introduction..... | 6 |
| 39 | 1 Scope | 7 |
| 40 | 2 Normative references | 7 |
| 41 | 3 Terms and definitions | 7 |
| 42 | 4 Conventions | 9 |
| 43 | 4.1 Reserved and unassigned values..... | 9 |
| 44 | 4.2 Byte ordering..... | 9 |
| 45 | 5 Assumptions | 9 |
| 46 | 5.1 Underlying Memory Mapping..... | 9 |
| 47 | 5.2 Multiple Instances | 9 |
| 48 | 5.3 Resets and Errors..... | 10 |
| 49 | 5.4 Notifications (Interrupts)..... | 10 |
| 50 | 5.5 Packet Sizes, Types, and Packet Flow..... | 10 |
| 51 | 5.6 Security | 11 |
| 52 | 6 Basic Architecture Concept..... | 11 |
| 53 | 7 MMBI Data Structures | 12 |
| 54 | 7.1 MMBI Capability Descriptor | 13 |
| 55 | 7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer..... | 15 |
| 56 | 7.2.1 Variable Packet Size Circular Buffer Descriptor | 15 |
| 57 | 7.2.2 Host Read-Write Structure..... | 17 |
| 58 | 7.2.3 Host Read-Only Structure..... | 18 |
| 59 | 8 Runtime Flows..... | 20 |
| 60 | 8.1 MMBI Interface Initialization and Reset | 20 |
| 61 | 8.1.1 Initialization of Descriptor Structures after Power Up | 20 |
| 62 | 8.1.2 Interface States and Graceful Reset..... | 21 |
| 63 | 8.1.3 Ungraceful Reset Considerations | 27 |
| 64 | 8.2 Calculation of Filled Space and Empty Space in Circular Buffer..... | 28 |
| 65 | 8.3 Device Readiness and Communication Pause | 28 |
| 66 | 8.4 Packet Transfer..... | 30 |
| 67 | 8.5 Interrupts (Optional)..... | 31 |
| 68 | 9 Multi-Protocol Packet Format..... | 31 |
| 69 | ANNEX A (informative) Notations | 33 |
| 70 | ANNEX B (informative) Change log..... | 34 |
| 71 | | |

72 **Figures**

| | | |
|----|---|----|
| 73 | Figure 1 – Multiple MMBI Instances..... | 10 |
| 74 | Figure 2 – MMBI Interface Concept Overview | 12 |
| 75 | Figure 3 – MMBI Data Structure Relationships..... | 13 |
| 76 | Figure 4 – MMBI Capability Descriptor Layout | 14 |
| 77 | Figure 5 – MMBI Interface States | 24 |
| 78 | Figure 6 – Sample MMBI Reset by Host..... | 25 |
| 79 | Figure 7 – Sample MMBI Reset by BMC | 27 |
| 80 | Figure 8 – Filled and Empty Space in Circular Buffers | 28 |
| 81 | Figure 9 – Sample MMBI Device Pause Sequences..... | 29 |
| 82 | | |

83 **Tables**

| | | |
|----|--|----|
| 84 | Table 1 – MMBI Capability Descriptor Structure (MMBI_Desc)..... | 14 |
| 85 | Table 2 – Buffer Type Dependent Descriptor for BUFT=0001b (VPSCB Descriptor) | 16 |
| 86 | Table 3 – MMBI Host Read-Write Structure (Host_RWS) | 18 |
| 87 | Table 4 – MMBI Host Read-Only Structure (Host_ROS)..... | 19 |
| 88 | Table 5 – MMBI Interface States | 22 |
| 89 | Table 6 – Multi-Protocol Packet Format..... | 32 |
| 90 | | |

91

Foreword

92 The *Memory-Mapped BMC Interface (MMBI) Specification* (DSP0282) was prepared by the DMTF PMCI
93 Working Group.

94 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
95 management and interoperability.

96 This version is the first version of this document. Future changes will be detailed in the change log in
97 ANNEX B.

98 Acknowledgments

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

119

Introduction

120 The Memory-Mapped BMC Interface (MMBI) specification defines the mechanisms facilitating
121 communication between platform management components, typically host software and a BMC
122 (baseboard management controller). Using the shared memory concept, this document defines a protocol
123 that allows packet exchanges between host software and BMC. The described memory mapping makes it
124 possible for both boot code (such as UEFI firmware), as well as OS-level software (such as OS kernel or
125 drivers) to establish efficient communication with BMC at bandwidth and latency limited by the underlying
126 memory mapping mechanisms.

127 1 Scope

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

133 2 Normative references

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

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

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

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

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

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

149 3 Terms and definitions

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

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

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

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

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

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

| | |
|-----|--|
| 166 | 3.1 |
| 167 | ACK |
| 168 | Acknowledge |
| 169 | 3.2 |
| 170 | B2H |
| 171 | BMC-to-Host |
| 172 | 3.3 |
| 173 | BAR |
| 174 | Base Address Register |
| 175 | 3.4 |
| 176 | CCT |
| 177 | Control Command Type |
| 178 | 3.5 |
| 179 | Destination Device |
| 180 | Device receiving the MCTP packet over MMBI |
| 181 | 3.6 |
| 182 | H2B |
| 183 | Host-to-BMC |
| 184 | 3.7 |
| 185 | MMBI |
| 186 | Memory-Mapped BMC Interface |
| 187 | 3.8 |
| 188 | MMIO |
| 189 | Memory-Mapped Input/Output |
| 190 | 3.9 |
| 191 | NACK |
| 192 | Not acknowledge |
| 193 | 3.10 |
| 194 | ROS |
| 195 | Read-Only Structure |
| 196 | 3.11 |
| 197 | RWS |
| 198 | Read-Write Structure |
| 199 | 3.12 |
| 200 | Source Device |
| 201 | Device sending the MCTP packet over MMBI |
| 202 | 3.13 |
| 203 | SPDM |
| 204 | Security Protocol and Data Model |

205 **3.14**
206 **VPSCB**
207 Variable Packet Size Circular Buffer

208 **4 Conventions**

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

210 **4.1 Reserved and unassigned values**

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

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

215 **4.2 Byte ordering**

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

218 **5 Assumptions**

219 **5.1 Underlying Memory Mapping**

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

224 Assumptions about the underlying layer are:

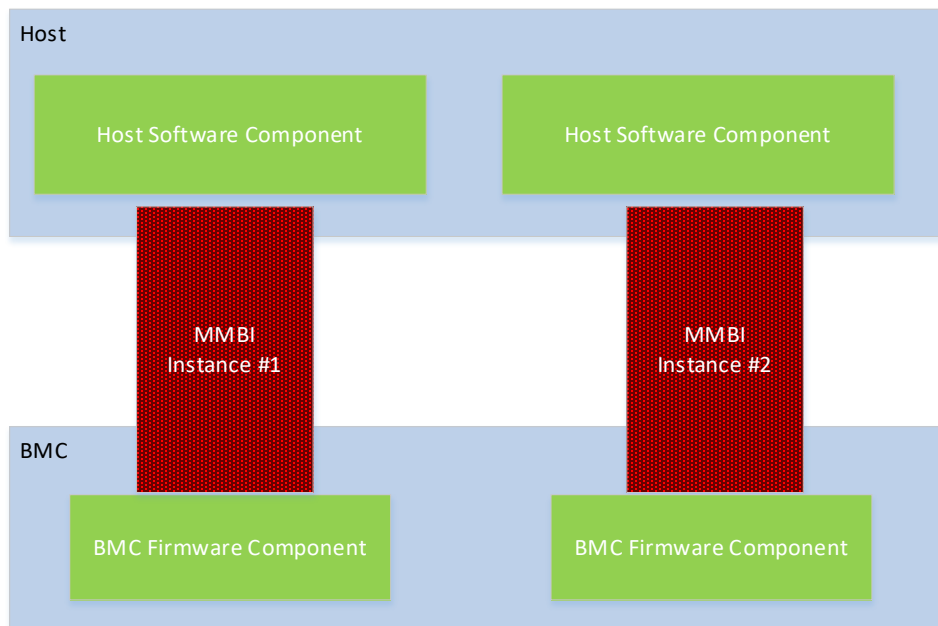
- 225 1) Memory mapping shall guarantee an error-free lossless channel.
- 226 2) The size of atomic operations is at least 4 bytes.
- 227 3) The order of operations must be preserved: writes must be visible to the other party in the order
228 they were executed by the sender; reads cannot be prefetched/cached;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

230 **5.2 Multiple Instances**

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

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

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



239

240

Figure 1 – Multiple MMBI Instances

241 **5.3 Resets and Errors**

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

247 **5.4 Notifications (Interrupts)**

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

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

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

255 **5.5 Packet Sizes, Types, and Packet Flow**

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

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

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

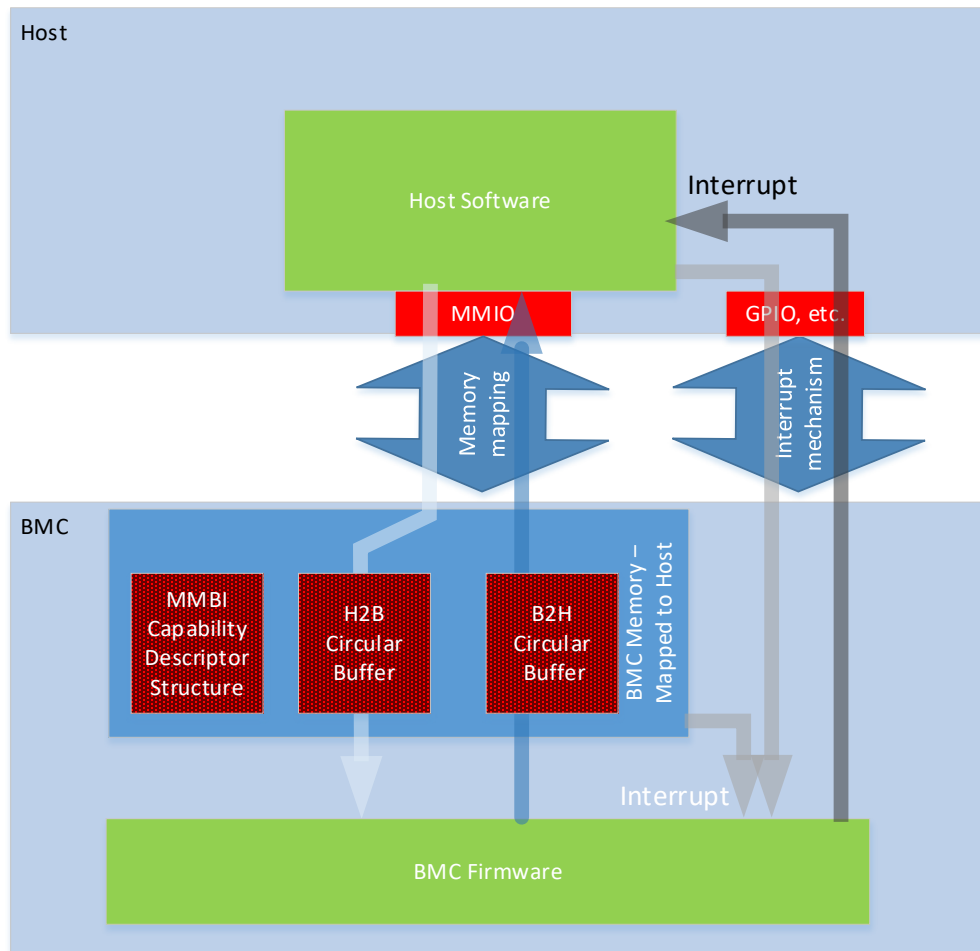
269 MMBI enables and defines discovery mechanisms to support the exchange of a variety of packet protocol
270 types, such as MCTP. Binding of these protocols to MMBI is defined in separate documents, such as
271 [Management Component Transport Protocol \(MCTP\) Memory-Mapped BMC Interface \(MMBI\) Transport](#)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

273 5.6 Security

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

279 6 Basic Architecture Concept

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



286

287

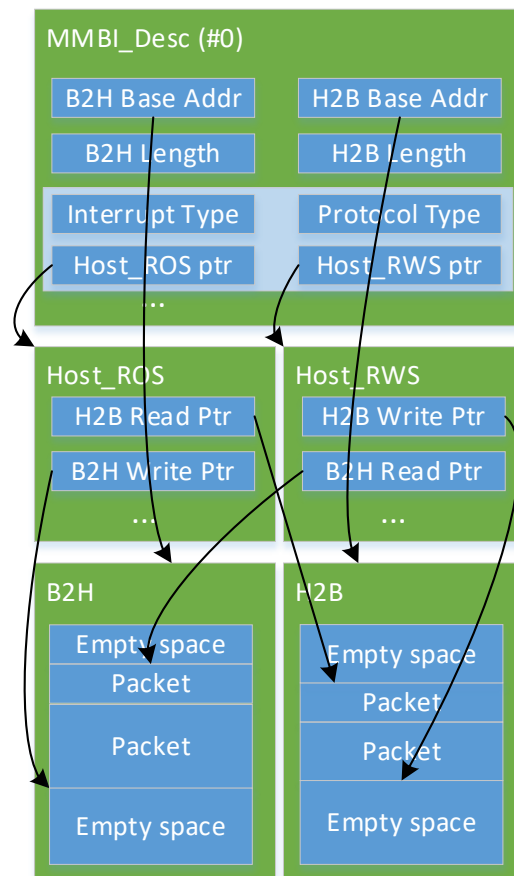
Figure 2 – MMBI Interface Concept Overview

288 7 MMBI Data Structures

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

- 290
- 291 • “BMC-to-Host” (B2H) region with substructure as follows:
 - 292 ○ *MMBI Capability Descriptor (MMBI_Desc* Structure)—see section 7.1 for details
 - 293 ○ *Host_ROS* (Host Read-Only Structure) —see section 7.2.3 for details
 - 294 ○ BMC-to-Host Circular buffer (B2H Circular buffer) —see section 8 for details
 - 295 • “Host-to-BMC” (H2B) region with substructure as follows:
 - 296 ○ *Host_RWS* (Host Read-Write Structure) —see section 7.2.2 for details
 - 297 ○ Host-to-BMC circular buffer (H2B Circular buffer) —see section 8 for details

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



301

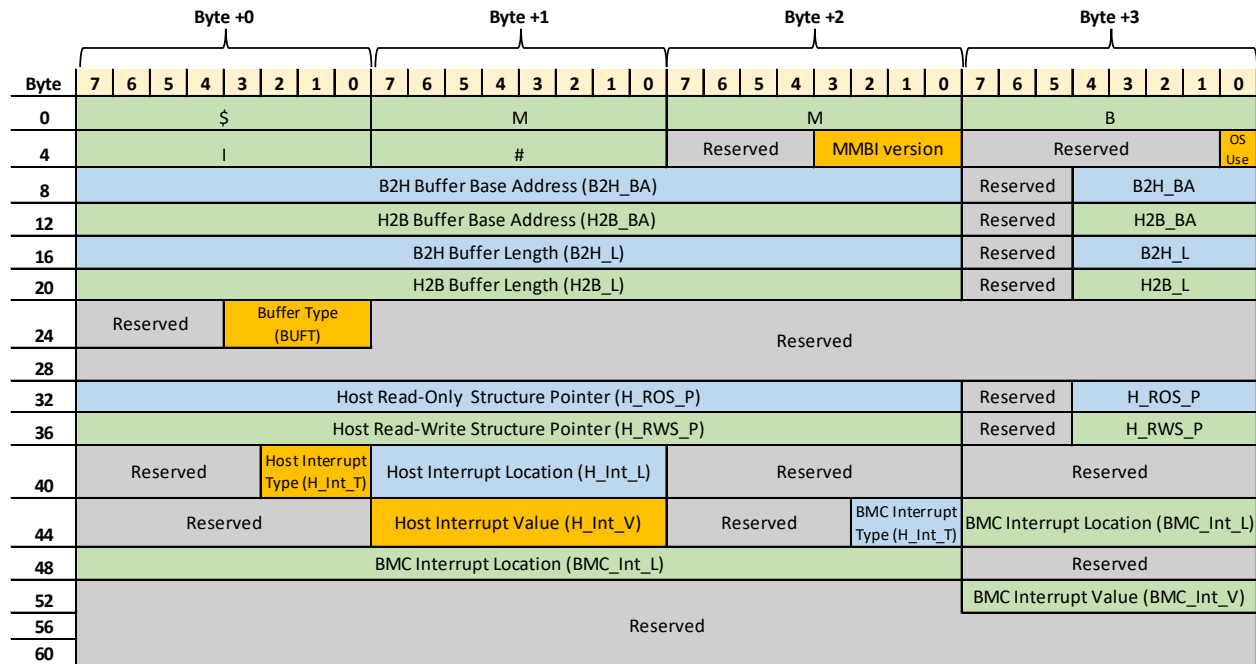
302

Figure 3 – MMBI Data Structure Relationships

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

306 **7.1 MMBI Capability Descriptor**

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



312

313

Figure 4 – MMBI Capability Descriptor Layout

314

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 |

315

316 **7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer**

317 This section describes data structures used when the communication between BMC and host SW
318 happens according to the VPSCB Buffer Type (BUFT=0001b).

319 **7.2.1 Variable Packet Size Circular Buffer Descriptor**

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

- 322
- The BMC updates this data structure during MMBI interface initialization.
 - Neither the BMC nor the host are allowed to update it at any other time.
- 323

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

| Byte(s) | Description |
|---------|---|
| 0:3 | [31:29] – Reserved |
| | <p>[28:0] Host Read-Only Structure Pointer (H_ROS_P)</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>[28:0] Host Read-Write Structure Pointer (H_RWS_P)</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>[2:0] Host Interrupt Type (H_Int_T)</p> <p>Defines how the BMC interrupts the host. This is an informative field from the host's perspective with the intention to keep the BMC 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>Host Interrupt Location (H_Int_L)</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>Host Interrupt Value (H_Int_V)</p> <p>If H_Int_T = 3: eSPI Virtual Wire data value Reserved otherwise</p> |

| Byte(s) | Description |
|---------|--|
| 14 | [7:3] – Reserved |
| | <p>[2:0] BMC Interrupt Type (BMC_Int_T)</p> <p>Defines how the BMC 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>BMC Interrupt Location (BMC_Int_L)</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>BMC Interrupt Value (BMC_Int_V)</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> |

325 **7.2.2 Host Read-Write Structure**

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

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

333

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

| Byte(s) | Description |
|---------|--|
| 0:3 | <p>[31:2] H2B Write Pointer (H2B_WP)</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 BMC 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 BMC, as indicated by the H2B_RP.</p> |
| | <p>[1] Host Interface Up (H_UP)</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 BMC.</p> |
| | <p>[0] Host Reset Request (H_RST)</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>[31:2] B2H Read Pointer (B2H_RP)</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 BMC 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>[1] Reserved</p> |
| | <p>[0] Host Ready (H_RDY)</p> <p>0 indicates that the host is performing some tasks that keep it busy, and so it may be unresponsive. However, the BMC 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> |

334 7.2.3 Host Read-Only Structure

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

339

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

| Byte(s) | Description |
|---------|---|
| 0:3 | <p>[31:2] B2H Write Pointer (B2H_WP)</p> <p>Bits [31:2] of the offset where the BMC 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 BMC shall advance the pointer once data is written to the Buffer to mark the next available offset.</p> <p>Note: BMC shall not overwrite the data not read by host, as indicated by the B2H_RP.</p> <hr/> <p>[1] BMC Interface Up (B_UP)</p> <p>1 indicates that the BMC side of the interface is up and running which means that the data structures are initialized and can be used</p> <hr/> <p>[0] BMC Reset Request (B_RST)</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>[31:2] H2B Read Pointer (H2B_RP)</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 BMC. Comparing this with the H2B write pointer will provide how much space is left to write.</p> <p>BMC shall only advance the pointer once the data available in H2B is read by the BMC.</p> <hr/> <p>[1] Reserved</p> <hr/> <p>[0] BMC Ready (B_RDY)</p> <p>0 indicates that the BMC 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 BMC is ready to exchange data (see section 8.1 for more information).</p> |

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

- 342 • H2B (Host-to-BMC buffer) is RW for the host and RO for the BMC.
- 343 • B2H (BMC-to-Host buffer) is RO for the host and RW for the BMC.

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

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

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

351 8 Runtime Flows

352 8.1 MMBI Interface Initialization and Reset

353 This section describes the steps to allow the BMC to complete the initialization of the data structures and
354 indicating when both sides of communication are ready to exchange data.

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

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

365 These flags are used to indicate the BMC's status as related to initialization and reset:

- 366 • BMC Interface Up (B_UP)
- 367 • BMC Reset Request (B_RST)

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

- 369 • Host Interface Up (H_UP)
- 370 • Host Reset Request (H_RST)

371 All these flags are used in combination to achieve the proper handshake mechanism between the host
372 and BMC during initialization or reset.

373 8.1.1 Initialization of Descriptor Structures after Power Up

374 The BMC must initialize the expected content of the MMBI data structures (see section 7) during power
375 up and make the shared memory available to the host. Initialization is expected to complete before the
376 host software accesses these structures so that the host can find the *MMBI Capability Descriptor*
377 (*MMBI_Desc*) using the MMBI signature bytes. MMBI structures and buffers must always remain
378 available in the shared memory when the host is using the MMBI interface.

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

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

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

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

393 **8.1.2 Interface States and Graceful Reset**

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

397 All the states are summarized in Table 5. The “Host Write Access” and “BMC Write Access” columns
398 define write-access restrictions to the data structures by host and BMC, respectively. There are no read
399 restrictions for the BMC and host. Note that the host is expected to re-read the data structure contents
400 after initialization is completed.

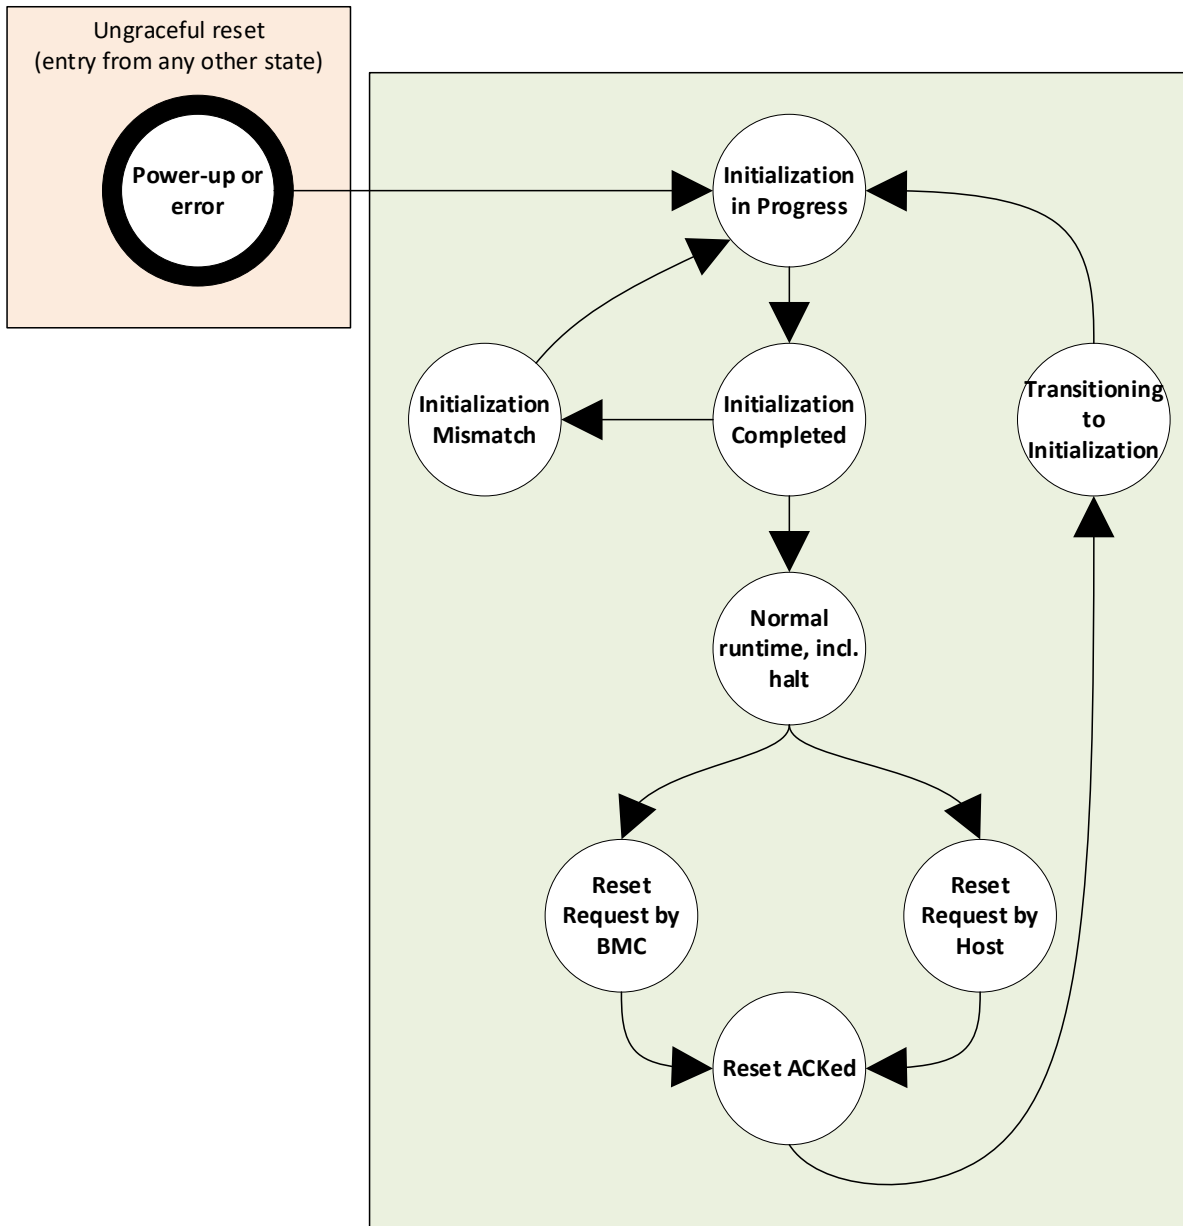
Table 5 – MMBI Interface States

| B_UP | B_RST | H_UP | H_RST | State Description | Host Write Access | BMC Write Access |
|------|-------|------|-------|--|---|--|
| 0 | 0 | 0 | 0 | <p>Initialization in Progress</p> <p>The BMC 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 | BMC allowed to write to any MMBI structures |
| 1 | 0 | 0 | 0 | <p>Initialization Completed</p> <p>The BMC 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 BMC is allowed to deposit packets into the circular buffer.</p> | Host allowed to write to MMBI structures as per section 7 | BMC allowed to write to MMBI structures as per section 7 |
| 1 | 0 | 1 | 0 | <p>Normal Runtime</p> <p>Both the BMC 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 | BMC allowed to write to MMBI structures as per section 7 |
| 1 | 1 | 1 | 0 | <p>Reset Request by BMC</p> <p>The BMC 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 | BMC allowed to write to MMBI structures as per section 7 |
| 1 | 0 | 1 | 1 | <p>Reset Request by Host</p> <p>The host is requesting reset—waiting for the BMC 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 BMC (B_UP = 1 and B_RST = 0 status).</p> <p>When the BMC 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 | BMC allowed to write to any MMBI structures |
| 1 | 1 | 1 | 1 | <p>Reset ACKed</p> <p>The host and BMC are ready to perform graceful interface reset. This is a transient state when the host is waiting for the BMC to complete the initialization. The host is not allowed to write to MMBI data structures. The BMC is expected to clear the B_UP and B_RST flags (in this order) and reinitialize all the data structures.</p> | Host not allowed to write to any MMBI structures | BMC allowed to write to any MMBI structures |

| B_UP | B_RST | H_UP | H_RST | State Description | Host Write Access | BMC Write Access |
|------|-------|------|-------|--|--|---|
| 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 | BMC allowed to write to any MMBI structures |
| 0 | 1 | 1 | 0 | Temporary Transition States These states may be observed during initialization when the BMC 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 BMC may initialize the interface or stop using MMBI and report a fatal error. | Host not allowed to write to any MMBI structures | BMC 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 | BMC allowed to write to any MMBI structures |
| 1 | 1 | 0 | 1 | Unexpected States | | |
| 1 | 1 | 0 | 0 | These states shall never happen: <ul style="list-style-type: none"> If the BMC reads this state, it indicates that the host does not follow MMBI protocol or some other corruption happened—the BMC should initialize the interface or it may stop using MMBI and report a fatal error, depending on system policy. 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. | | |
| 0 | 0 | 1 | 0 | | | |
| 0 | 0 | 1 | 1 | | | |

402 The expected state transitions are presented in Figure 5:

403



404

405

Figure 5 – MMBI Interface States

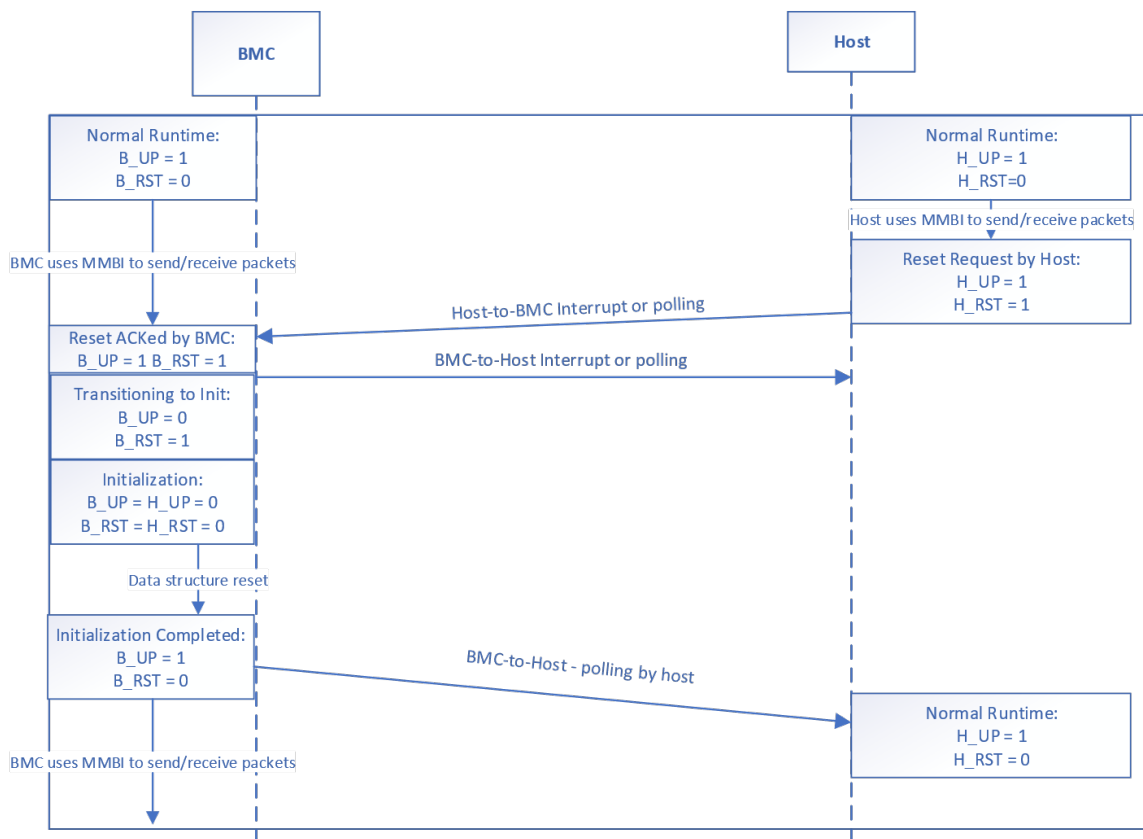
406 The host shall check the MMBI Interface state before writing any new data to the H2B buffer (as
 407 described in Table 5, the host is only allowed to transfer new data in the Normal Runtime state, i.e.,
 408 B_UP=1 & B_RST=0 & H_UP=1 & H_RST=0). Similarly, the BMC shall check the status before writing
 409 any new data to the B2H buffer. These status flags are conveniently located in the B2H_WP or H2B_WP
 410 bytes which the host or BMC, respectively, read anyway during any use of the circular buffers.

411 **8.1.2.1 Host Initiating Graceful Reset Sequence**

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

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

426 Figure 6 presents a sample flow:



427

428

Figure 6 – Sample MMBI Reset by Host

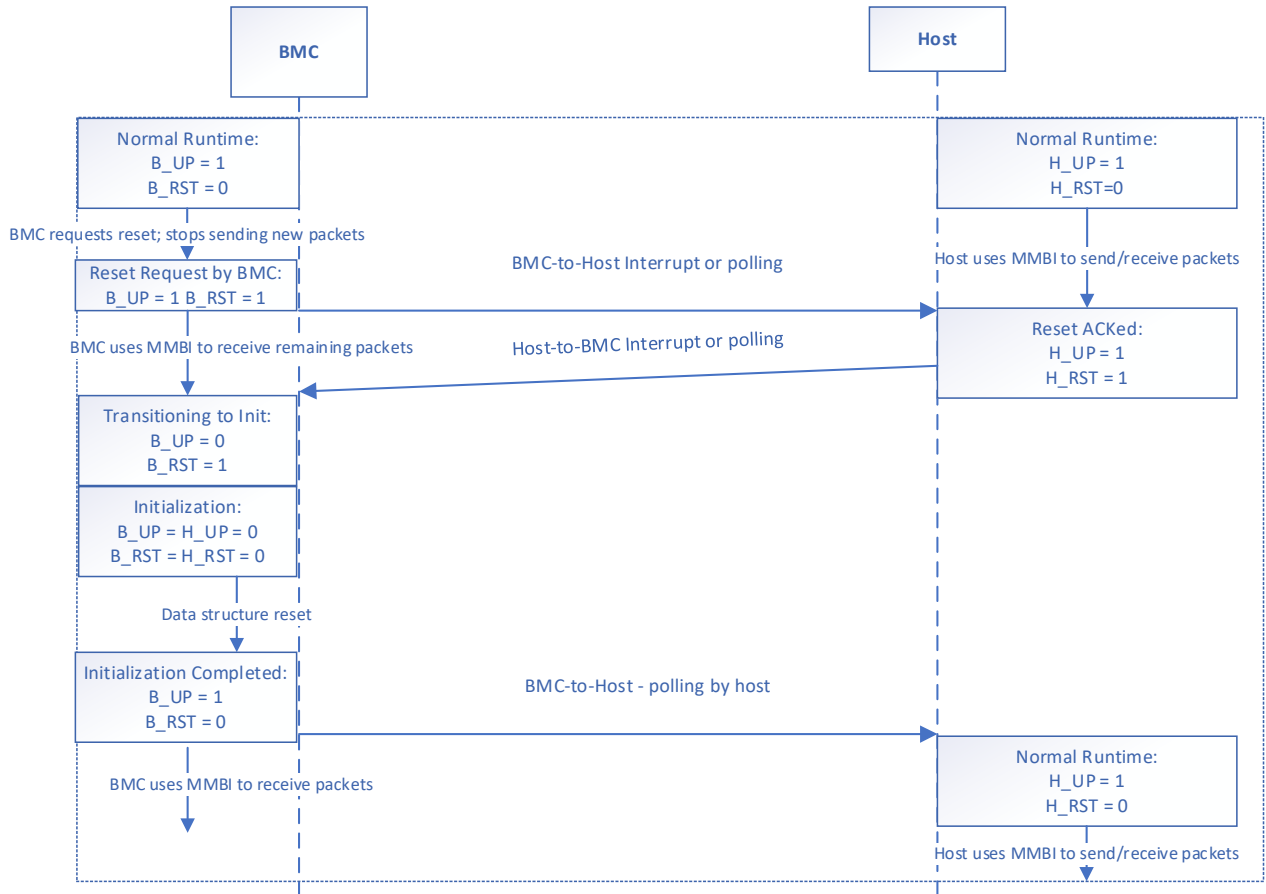
429 **8.1.2.2 BMC Initiating Graceful Reset Sequence**

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

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

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

448 Figure 7 presents a sample flow.



449

450

Figure 7 – Sample MMBI Reset by BMC

451 **8.1.3 Ungraceful Reset Considerations**

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

- 454
- 455
- 456
- 457
- 458
- 459 1. In the case of a BMC FW-only reset (HW continues to work, memory content, including
460 buffers stay intact in shared memory and accesses are still handled by HW): the host will still
461 see the MMBI in the normal state and write to MMBI Circular buffers to deposit or read data
462 as long as there is any space available in the buffers. In this situation, host may timeout
463 waiting for a response but this is handled by higher layers above MMBI.
 - 464 2. BMC HW reset (buffers are wiped and MMIO mechanisms are broken): the host will see
465 errors on reads/writes and must handle them as per host-specific mechanisms. Additionally,
466 MMBI encoding of status in B_UP, B_RST, H_UP, & H_RST is such that all-zeros or all-ones
467 are recognized as transient states (see Table 5). So, even if there would be no other
468 mechanisms in the system, the host would still recognize this as an error and would have to
469 wait for reinitialization by the BMC (the host is not allowed to write to the buffers in the
transient state, i.e., until the data structures are reinitialized by BMC 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 BMC memory, so the BMC can still read the data from
the buffers. Assuming the BMC understands the host's status via other mechanisms, the
BMC can take informed decisions about how to respond to such situations.

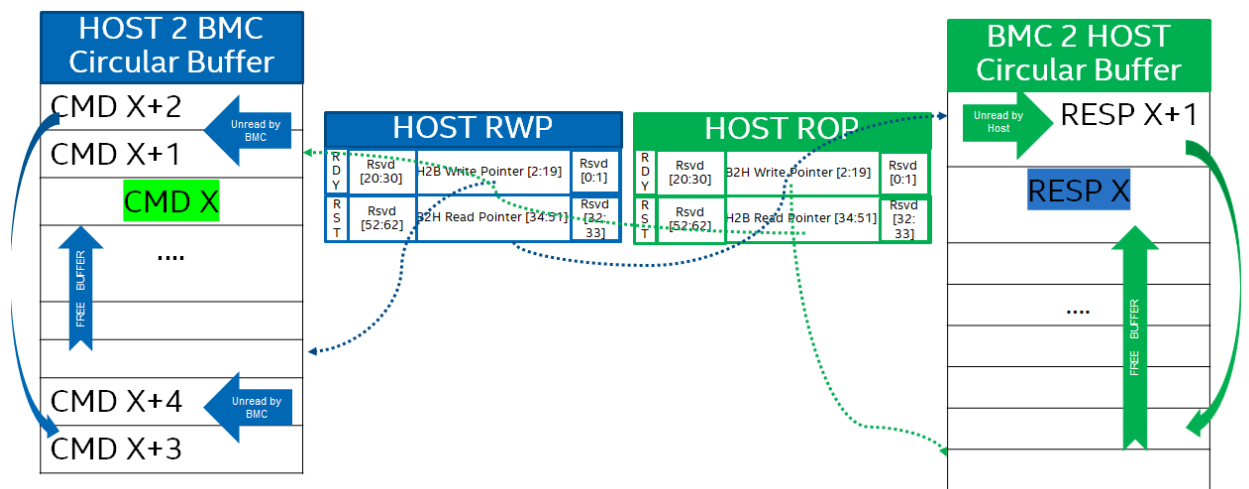
470 In all the above cases, MMBI data structures can be reinitialized after the reset to allow a clean restart.

471 **8.2 Calculation of Filled Space and Empty Space in Circular Buffer**

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

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

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



484

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

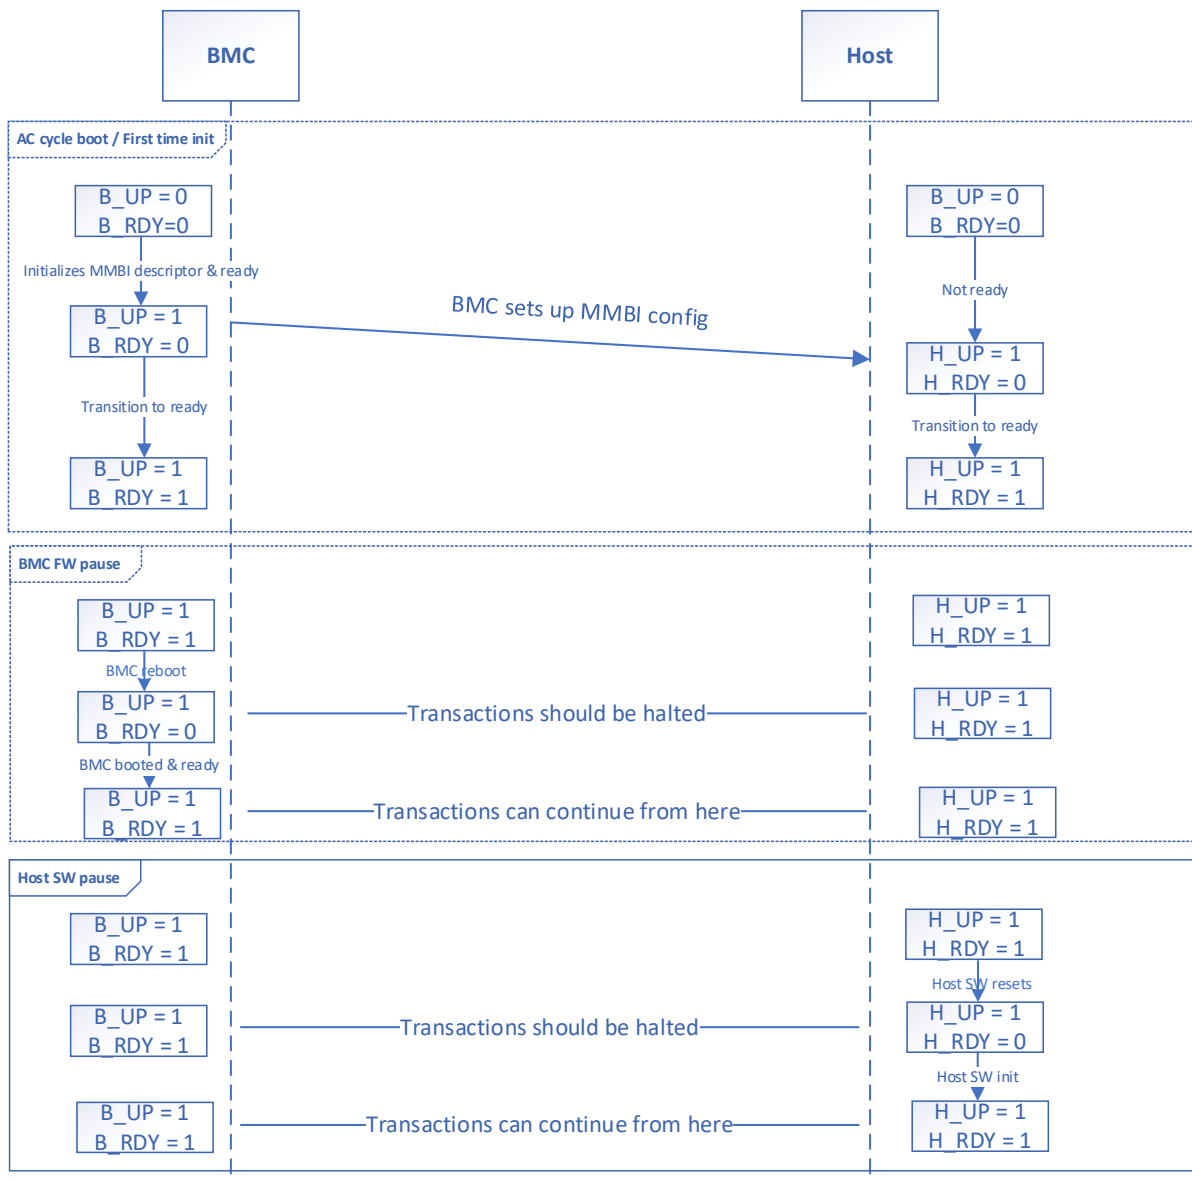
486 **8.3 Device Readiness and Communication Pause**

487 In addition to the reinitialization or reset states, the MMBI interface also uses the H_RDY and B_RDY
 488 flags to indicate the device's readiness to consume incoming packets and handle them. When the host or
 489 BMC are ready to receive and handle packets, they set the B_RDY or H_RDY flags, respectively. If a
 490 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
 491 up but the target device is not ready to consume and handle new packets. When the interface is up, it
 492 means that the data structures are ready to accept new packets so the sender can:

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

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

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



502

503

Figure 9 – Sample MMBI Device Pause Sequences

504 8.4 Packet Transfer

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

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

542 8.5 Interrupts (Optional)

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

- 545 • A packet has just been written to the circular buffer.
- 546 • A packet has just been read from the circular buffer.
- 547 • The host or BMC has initiated an interface reset sequence.
- 548 • The host or BMC has completed its portion of the interface reset sequence and normal operation
549 can begin.

550 An interrupt handler shall:

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

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

559 9 Multi-Protocol Packet Format

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

564

Table 6 – Multi-Protocol Packet Format

| Byte(s) | Description | | | | | | |
|---------|--|---------|-------------|-----|--|-------|-------------------------------|
| 0:2 | <p>[23:2] Packet Length (PKT_LEN)</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>[1:0] Packet padding (PKT_PAD)</p> <p>Number of padding bytes</p> | | | | | | |
| 3 | <p>[7:4] – Reserved</p> | | | | | | |
| | <p>[3:0] – Packet type (PKT_TYPE)</p> <p>Defines the format of the remaining bytes:</p> <p>0100b – MCTP over MMBI (see Management Component Transport Protocol (MCTP) Memory-Mapped BMC Interface (MMBI) Transport Binding Specification)</p> <p>0101b – Vendor defined content as defined below</p> <p>Other values are reserved.</p> | | | | | | |
| 4:N-1 | <p>Protocol type specific fields</p> <p>This field depends on the PKT_TYPE value:</p> <p>If PKT_TYPE = MCTP = 0100b: format follows MCTP over MMBI (see Management Component Transport Protocol (MCTP) Memory-Mapped BMC Interface (MMBI) Transport Binding Specification)</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 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) | <p>Padding (PAD) – optional</p> <p>Padding bytes as defined in PKT_PAD field.</p> <p>Note: padding is added to ensure packets are 4-byte aligned</p> | | | | | | |

565

ANNEX A (informative)

Notations

566
567
568
569
570

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

- 572 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets
573 starting from byte two and continuing to and including byte N. The lowest offset is on
574 the left; the highest is on the right.
- 575 • (6) Parentheses around a single number can be used in packet field descriptions to
576 indicate a byte field that may be present or absent.
- 577 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range
578 may be present or absent. The lowest offset is on the left; the highest is on the right.
- 579 • [PCle](#) Underlined blue text is typically used to indicate a reference to a document or
580 specification called out in clause 2, "Normative References" or to items hyperlinked
581 within the document.
- 582 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets
583 are given as zero-based values (that is, the least significant bit offset = 0).
- 584 • [7:5] A range of bit offsets. The most significant bit is on the left, the least significant bit is
585 on the right.
- 586 • 1b A number consisting of 0s and 1s followed by a lowercase "b" indicates that the
587 number is in binary format.
- 588 • 0x12A A leading "0x" indicates that the number is in hexadecimal format.

589
590
591
592
593

ANNEX B (informative)

Change log

| Version | Date | Description |
|---------|------------|------------------|
| 1.0.0 | 2023-08-25 | Initial release. |

594