

1



2

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

5 **Version: 1.1.0**

6 **Document Identifier: DSP0282**

7 **Date: 2025-12-17**

8 **Supersedes: 1.0.2**

9 **Document Class: Normative**

10 **Document Status: Published**

11 **Document Language: en-US**

12 Copyright Notice

13 Copyright © 2023–2025 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29 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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 PCI-SIG®, PCI Express®, and PCIe® are registered trademarks or service marks of PCI-SIG. All other
36 marks and brands are the property of their respective owners.

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

CONTENTS

39 Foreword 6

40 Introduction 7

41 1 Scope 8

42 2 Normative references 8

43 3 Terms and definitions 8

44 4 Conventions 10

45 4.1 Reserved and unassigned values 10

46 4.2 Byte ordering 10

47 5 Assumptions 10

48 5.1 Underlying Memory Mapping 10

49 5.2 Multiple Instances 10

50 5.3 Resets and Errors 11

51 5.4 Notifications (Interrupts) 11

52 5.5 Packet Sizes, Types, and Packet Flow 12

53 5.6 Security 12

54 6 Basic Architecture Concept 12

55 7 MMBI Data Structures 13

56 7.1 MMBI Capability Descriptor 14

57 7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer 17

58 7.2.1 Variable Packet Size Circular Buffer Descriptor 18

59 7.2.2 Host Read-Write Structure 19

60 7.2.3 Host Read-Only Structure 20

61 8 Runtime Flows 22

62 8.1 MMBI Interface Initialization and Reset 22

63 8.1.1 Initialization of Descriptor Structures after Power Up 22

64 8.1.2 Interface States and Graceful Reset 23

65 8.1.3 Ungraceful Reset Considerations 29

66 8.2 Calculation of Filled Space and Empty Space in Circular Buffer 30

67 8.3 Device Readiness and Communication Pause 30

68 8.4 Packet Transfer 32

69 8.5 Interrupts (Optional) 33

70 9 Multi-Protocol Packet Format 33

71 ANNEX A (informative) Notations 35

72 ANNEX B (informative) Change log 36

73

74 **Figures**

75	Figure 1 – Multiple MMBI Instances.....	11
76	Figure 2 – MMBI Interface Concept Overview.....	13
77	Figure 3 – MMBI Data Structure Relationships.....	14
78	Figure 4 – MMBI Capability Descriptor Layout with Multi-Channels.....	15
79	Figure 5 – MMBI Interface States	26
80	Figure 6 – Sample MMBI Reset by Host.....	27
81	Figure 7 – Sample MMBI Reset Flow by (B)MC.....	29
82	Figure 8 – Filled and Empty Space in Circular Buffers	30
83	Figure 9 – Sample MMBI Device Pause Sequences.....	31
84		

85 **Tables**

86 Table 1 – MMBI Capability Descriptor Structure (MMBI_Desc) with Multi-Channels 15

87 Table 2 – Buffer Type Dependent Descriptor for BUFT=0001b (VPSCB Descriptor) 18

88 Table 3 – MMBI Host Read-Write Structure (Host_RWS) 20

89 Table 4 – MMBI Host Read-Only Structure (Host_ROS)..... 21

90 Table 5 – MMBI Interface States 24

91 Table 6 – Multi-Protocol Packet Format..... 34

92

93

Foreword

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

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

98 This version of the specification supersedes version 1.0.2. Changes are detailed in the change log in
99 ANNEX B.

100 Acknowledgments

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

102 Editors:

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

107 Contributors:

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

124

Introduction

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

135 1 Scope

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

141 2 Normative references

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

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

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

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

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

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

157 PCI-SIG, PCI Express® Base Specification Revision 6.2, January 25, 2024
158 <https://www.pcisig.com/specifications/>

159 3 Terms and definitions

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

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

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

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

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

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

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

215 **3.14**

216 **Source Device**

217 Device sending the MCTP packet over MMBI

218 **3.15**

219 **SPDM**

220 Security Protocol and Data Model

221 **3.16**

222 **VPSCB**

223 Variable Packet Size Circular Buffer

224 **4 Conventions**

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

226 **4.1 Reserved and unassigned values**

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

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

231 **4.2 Byte ordering**

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

234 **5 Assumptions**

235 **5.1 Underlying Memory Mapping**

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

240 Assumptions about the underlying layer are:

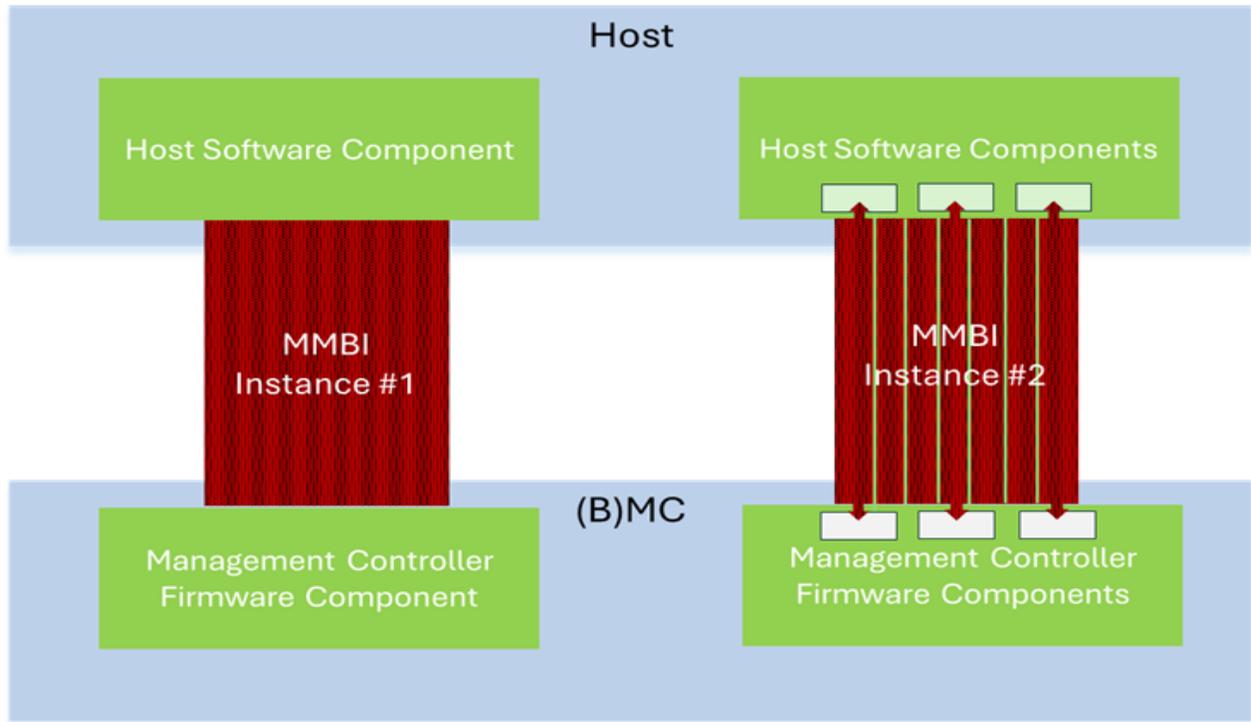
- 241 1) Memory mapping shall guarantee an error-free lossless channel.
242 2) The size of atomic operations is at least 4 bytes.
243 3) The order of operations must be preserved: writes must be visible to the other party in the order
244 they were executed by the sender; reads cannot be prefetched/cached;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

246 **5.2 Multiple Instances**

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

251 Multiple components in the system, e.g., multiple host tenant / software agents communicating to a
 252 (B)MC, can be supported using a plurality of MMBI interfaces (each being an independent instance of the
 253 interface), located in different memory locations or using channels within a single MMBI instance. Such
 254 MMBI instances shall operate independently, as shown in Figure 1. In the figure, Instance #1 is Single
 255 Channel, and Instance #2 is Multi Channel.

256



257

258

Figure 1 – Multiple MMBI Instances

259 **5.3 Resets and Errors**

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

265 **5.4 Notifications (Interrupts)**

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

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

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

273 5.5 Packet Sizes, Types, and Packet Flow

274 MMBI allows variable packet sizes, with the maximum size dependent on the underlying physical layer's
275 memory mapping capabilities. MMBI provides a discovery method allowing the communicating parties to
276 define and discover the circular buffer sizes, which limit the maximum packet sizes that can be
277 transmitted (fragmentation/reassembly is not supported by this version of MMBI protocol). The upper
278 layers must adhere to the discovered limits and, if necessary, handle fragmentation/reassembly
279 accordingly.

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

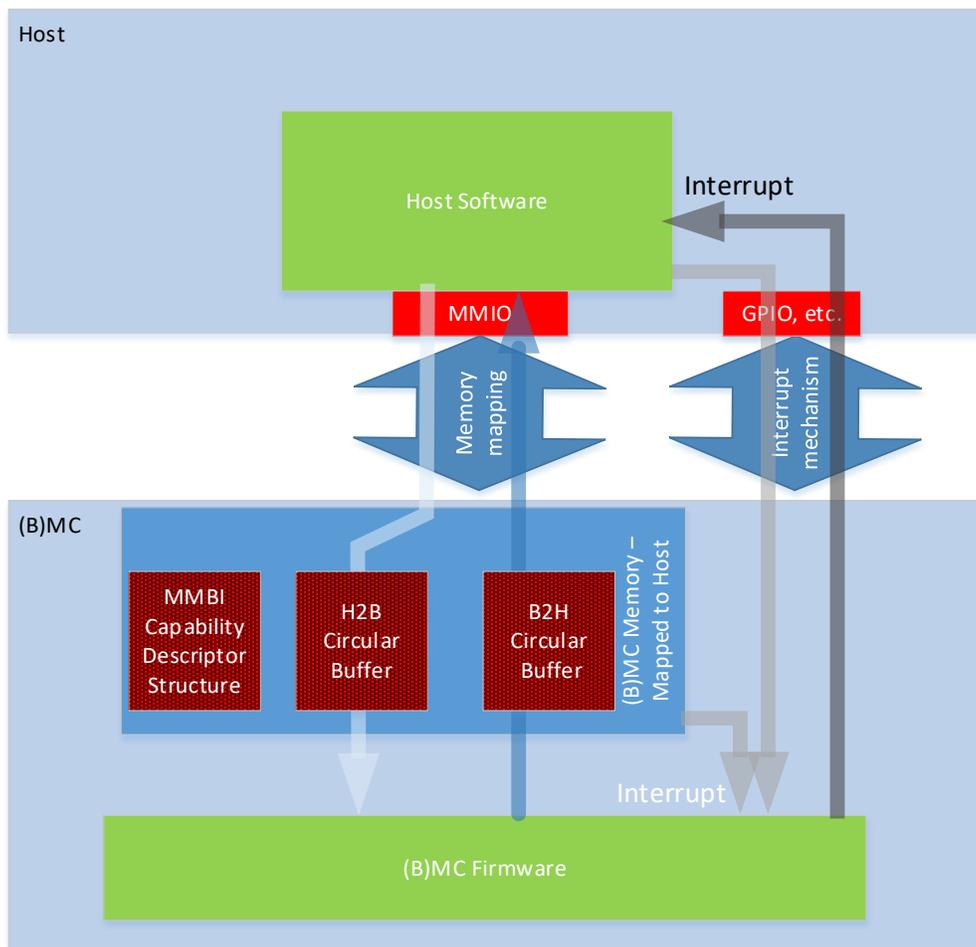
287 MMBI enables and defines discovery mechanisms to support the exchange of a variety of packet protocol
288 types, such as MCTP. Binding of these protocols to MMBI is defined in separate documents, such as
289 [Management Component Transport Protocol \(MCTP\) Memory-Mapped Buffer Interface \(MMBI\)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

291 5.6 Security

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

297 6 Basic Architecture Concept

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



304

305

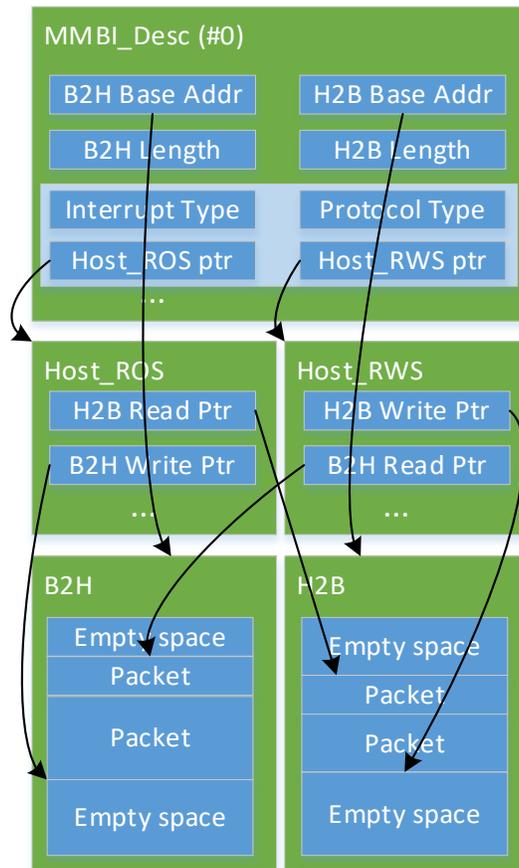
Figure 2 – MMBI Interface Concept Overview

306 7 MMBI Data Structures

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

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

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



319

320

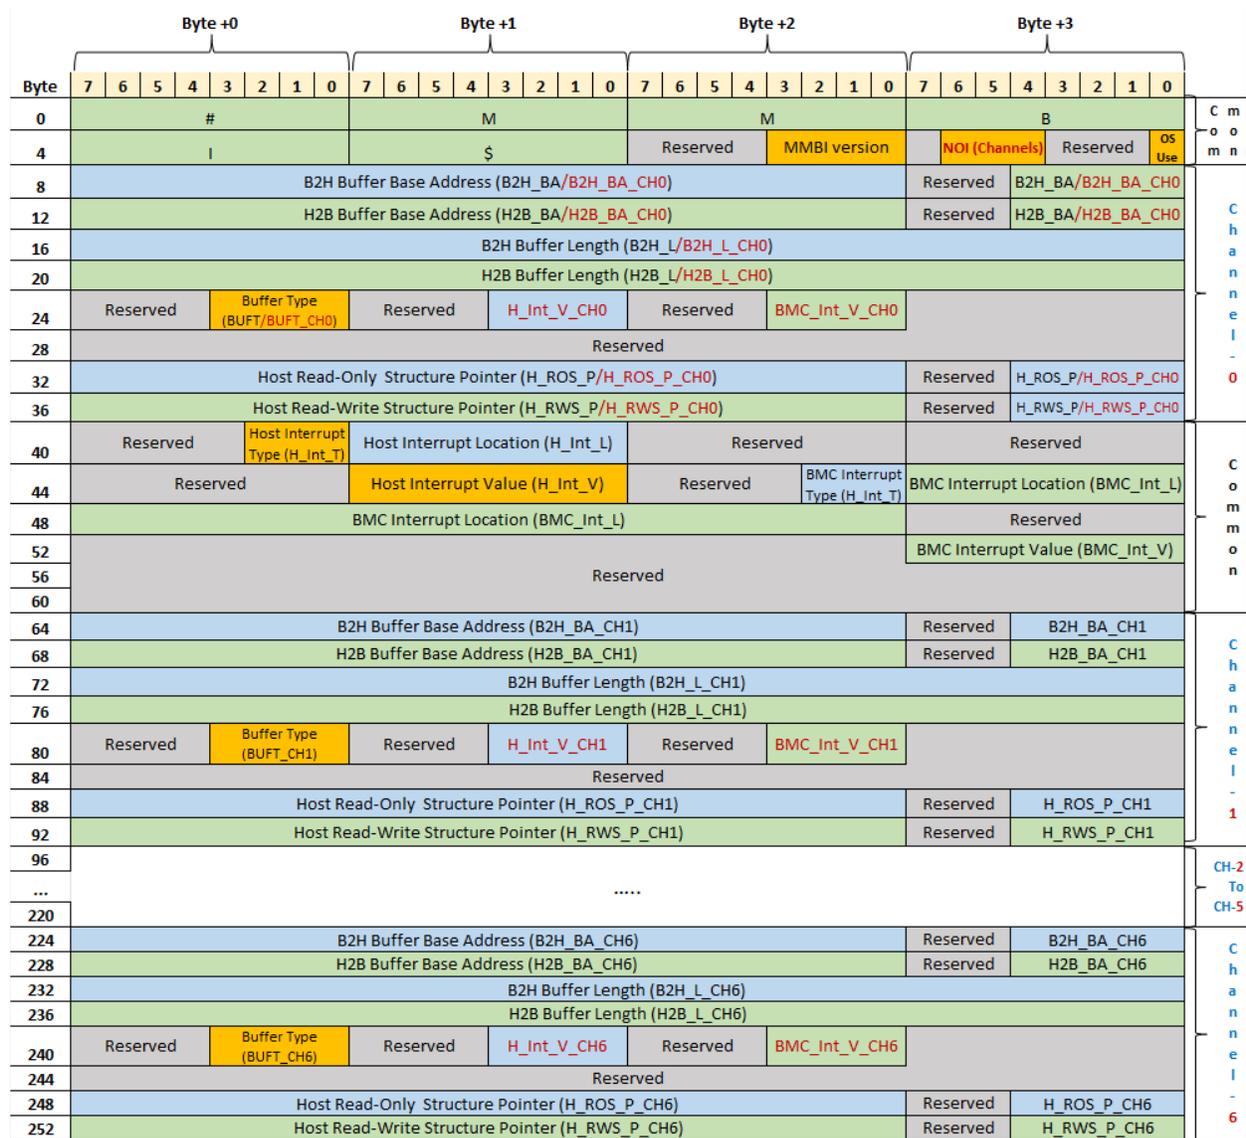
Figure 3 – MMBI Data Structure Relationships

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

324 **7.1 MMBI Capability Descriptor**

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

330



331

332

Figure 4 – MMBI Capability Descriptor Layout with Multi-Channels

333

Table 1 – MMBI Capability Descriptor Structure (MMBI_Desc) with Multi-Channels

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

Byte(s)	Description
	<p>[3:0] MMBI version</p> <p>0001b – Implementations of MMBI described in this document shall indicate version 1.0 of MMBI (Legacy Implementation).</p> <p>0010b – Implementations of MMBI described in this document shall indicate version 1.1 of MMBI.</p>
7	<p>[7] Reserved</p> <p>[6:4] NOI Number Of Instances (NOI) / Number of Memory channels created on the BAR/MMIO region. The supported channels (0–6).</p> <p>000b – Channel-0 active, and the header length is fixed at 64 bytes</p> <p>001b – Two Channels (CH0 – CH1)</p> <p>010b – Three Channels (CH0 – CH2)</p> <p>...</p> <p>110b – Seven Channels (CH0 – CH6)</p> <p>Other values are reserved</p> <hr/> <p>[3:1] Reserved</p> <p>[0] OS Use Indicates if this MMBI interface is intended for OS use:</p> <p>0b – OS should not use this MMBI interface as it is managed by other host software components (UEFI BIOS, ACPI ASL code, etc.).</p> <p>1b – This MMBI interface is intended for OS use</p>
8:11	<p>[31:29] – Reserved</p> <hr/> <p>[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</p>
12:15	<p>[31:29] – Reserved</p> <hr/> <p>[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</p>
16:19	<p>B2H Buffer Length (B2H_L) The size of the B2H buffer (can represent up to 4GB)</p>
20:23	<p>H2B Buffer Length (H2B_L) The size of the B2H buffer (can represent up to 4GB)</p>
24	<p>[7:4] Reserved</p>

Byte(s)	Description
	<p>[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.</p>
25	<p>[7:4] Reserved</p> <p>[3:0] H_Int_V_CH0 The Host Interrupt bit value associated with channel (CH0) in H_Int_V 0000b – No interrupt set 0001b – Interrupt set channel 0 Other values are reserved</p>
26	<p>[7:4] Reserved</p> <p>[3:0] BMC_Int_V_CH0 The BMC Interrupt bit value associated with channel (CH0) in BMC_Int_V 0000b – No interrupt set 0001b – Interrupt set channel 0 Other values are reserved</p>
27:31	Reserved
32:52	<p>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</p>
56:63	Reserved
64:95	<p>Legacy Mode Operation: If NOI == 0, the MMBI device operates in legacy mode with only Channel-0 active, and the header length is fixed at 64 bytes.</p> <p>Channel-1 Definition: Identical to the structure defined in Byte(s) 8 to 39</p> <ul style="list-style-type: none"> - Each additional channel beyond Channel-1 (CH1) increases the MMBI header length by 32 bytes. - The header can be extended to support up to 7 channels (CH0 to CH6). - With 7 channels (CH0 to CH6), the total MMBI header length reaches 256 bytes.

334

335 **7.2 MMBI Circular Buffers—Variable Packet Size Circular Buffer**

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

338 **7.2.1 Variable Packet Size Circular Buffer Descriptor**

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

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

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

Byte(s)	Description
0:3	[31:29] – Reserved
	[28:0] Host Read-Only Structure Pointer (H_ROS_P/H_ROS_P_CH0) Points to the <i>Host_ROS</i> structure. The base address is expressed in 8-byte units as the offset relative to beginning of the descriptor
4:7	[31:29] – Reserved
	[28:0] Host Read-Write Structure Pointer (H_RWS_P/H_RWS_P_CH0) Points to the <i>Host_RWS</i> structure. The base address is expressed in 8-byte units as the offset relative to beginning of the descriptor
8	[7:3] – Reserved
	[2:0] Host Interrupt Type (H_Int_T) Defines how the (B)MC interrupts the host. This is an informative field from the host's perspective with the intention to keep the (B)MC and host in sync. 0 – no interrupt / polling 1 – PCIe interrupt (bus specific) 2 – physical pin (GPIO) 3 – eSPI Virtual Wire Other values are reserved
9	Host Interrupt Location (H_Int_L) If H_Int_T = 0: reserved If H_Int_T = 1: for PCIe, indicates the PCIe interrupt message number If H_Int_T = 2: pin number If H_Int_T = 3: eSPI Virtual Wire Index number Reserved otherwise
10:12	Reserved
13	Host Interrupt Value (H_Int_V) If H_Int_T = 3: eSPI Virtual Wire data value If H_Int_T = 1: for PCIe, H_Int_V indicates the PCIe interrupt number bits corresponding to channels (CH0 - CH6) and the value to be written at the given address to trigger an interrupt for the (B)MC. Reserved otherwise

Byte(s)	Description
14	[7:3] – Reserved
	<p>[2:0] (B)MC Interrupt Type (BMC_Int_T)</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)MC 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>(B)MC 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>If BMC_Int_T = 2: Inband interrupt, BMC_Int_V indicates the PCIe interrupt bits corresponding to channels (CH0 - CH6) and the value to be written at the given address to trigger an interrupt for the Host.</p> <p>Otherwise reserved</p>

344 **7.2.2 Host Read-Write Structure**

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

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

352

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 (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>[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 (B)MC.</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 (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>[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 (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>

353 7.2.3 Host Read-Only Structure

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

358

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 (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>[1] (B)MC Interface Up (B_UP)</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>[0] (B)MC 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 (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>[1] Reserved</p>
	<p>[0] (B)MC Ready (B_RDY)</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>

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

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

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

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

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

370 8 Runtime Flows

371 8.1 MMBI Interface Initialization and Reset

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

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

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

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

- 385 • (B)MC Interface Up (B_UP)
- 386 • (B)MC Reset Request (B_RST)

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

- 388 • Host Interface Up (H_UP)
- 389 • Host Reset Request (H_RST)

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

392 8.1.1 Initialization of Descriptor Structures after Power Up

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

398 If the MMBI is made available via a memory mapped range of a PCIe function, then the *MMBI_Desc* is at
399 offset 0 of a PCIe function BARs, and the function PCIe Base Class, Subclass, and Programming shall be
400 {0xC, 0xC, 0x0}. There can be at most one *MMBI_Desc* per BAR. During the initial accesses after the
401 host's power up or reset, the host's software is expected to verify if the content of the MMBI version and
402 MMBI signature are as expected. If the above requirements are met, the host is expected to check the
403 interface state.

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

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

415 **8.1.2 Interface States and Graceful Reset**

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

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

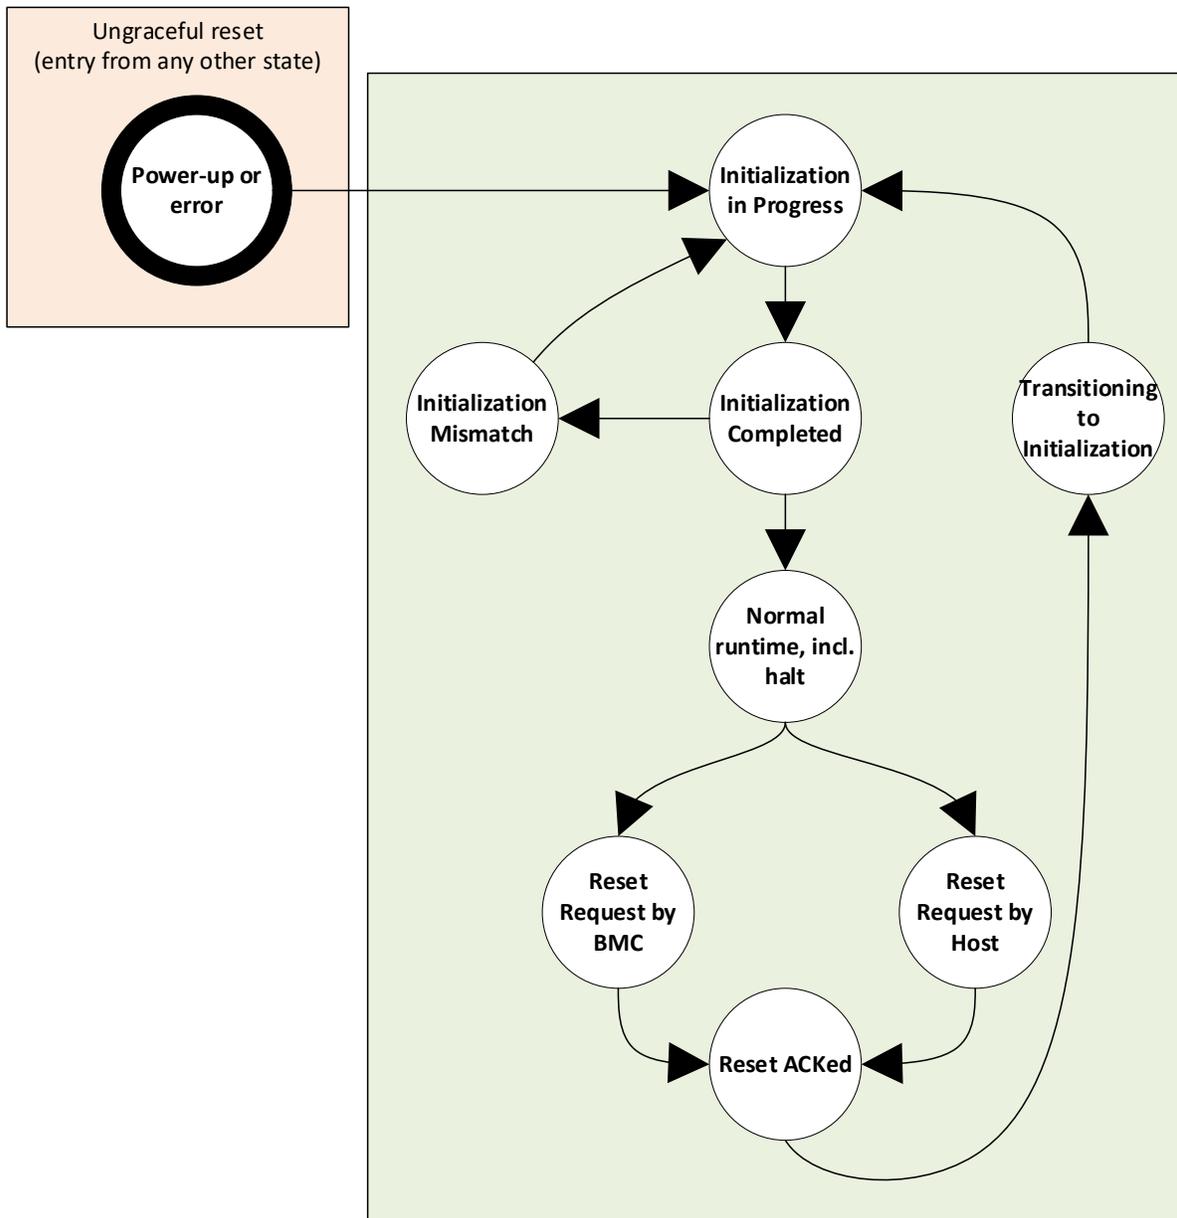
Table 5 – MMBI Interface States

B_UP	B_RST	H_UP	H_RST	State Description	Host Write Access	(B)MC Write Access
0	0	0	0	<p>Initialization in Progress</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>Initialization Completed</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>Normal Runtime</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>Reset Request by (B)MC</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>Reset Request by Host</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	<p>Reset ACKed</p> <p>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.</p>	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
0	1	1	1	<p>Transitioning to Initialization</p> <p>Transient state after the “Reset ACKed” state. The host is not allowed to write to MMBI data structures.</p>	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
0	1	1	0	<p>Temporary Transition States</p> <p>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).</p> <p>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.</p>	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	<p>Initialization Mismatch</p> <p>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.</p>	Host not allowed to write to any MMBI structures	(B)MC allowed to write to any MMBI structures
1	1	0	1	<p>Unexpected States</p> <p>These states shall never happen:</p> <ul style="list-style-type: none"> 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. 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. 		
1	1	0	0			
0	0	1	0			
0	0	1	1			

424 The expected state transitions are presented in Figure 5:

425



426

427

Figure 5 – MMBI Interface States

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

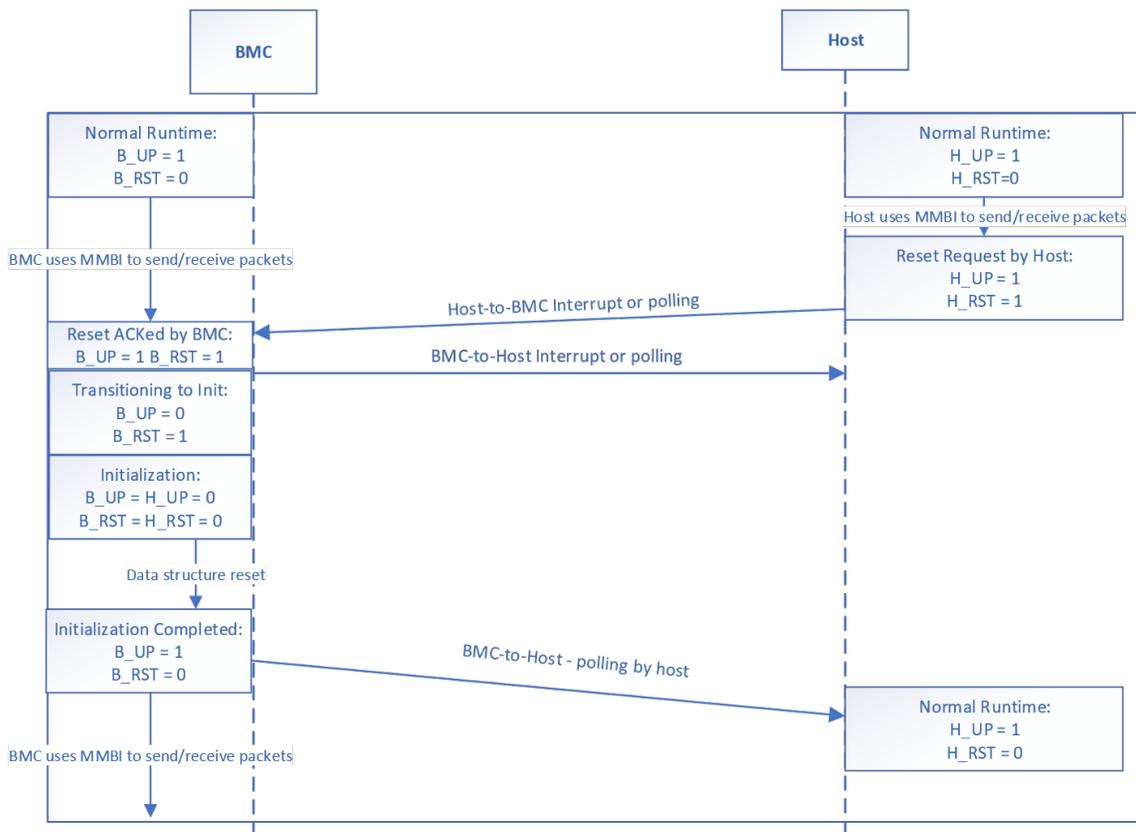
431 any new data to the B2H buffer. These status flags are conveniently located in the B2H_WP or H2B_WP
 432 bytes which the host or (B)MC, respectively, read anyway during any use of the circular buffers.

433 **8.1.2.1 Host Initiating Graceful Reset Sequence**

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

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

448 Figure 6 presents a sample flow:



449

450

Figure 6 – Sample MMBI Reset by Host

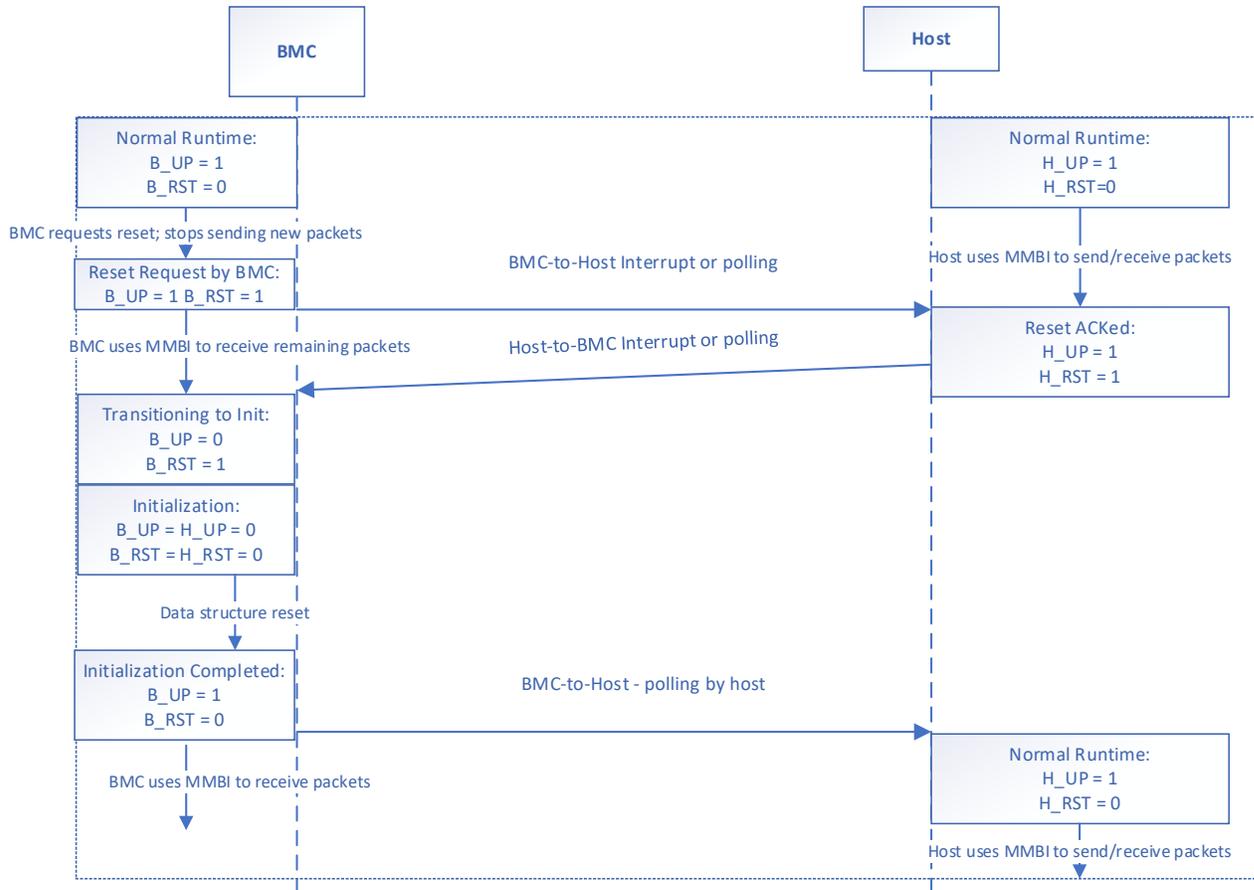
451 8.1.2.2 (B)MC Initiating Graceful Reset Sequence

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

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

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

470 Figure 7 presents a sample flow.



471

472

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

473 8.1.3 Ungraceful Reset Considerations

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

- 476
- 477
- 478
- 479
- 480
1. In the case of a (B)MC FW-only reset (HW continues to work, memory content, including buffers stay intact in shared memory and accesses are still handled by HW): the host will still see the MMBI in the normal state and write to MMBI Circular buffers to deposit or read data as long as there is any space available in the buffers. In this situation, host may timeout waiting for a response but this is handled by higher layers above MMBI.
 2. (B)MC HW reset (buffers are wiped and MMIO mechanisms are broken): the host will see errors on reads/writes and must handle them as per host-specific mechanisms. Additionally, MMBI encoding of status in B_UP, B_RST, H_UP, & H_RST is such that all-zeros or all-ones are recognized as transient states (see Table 5). So, even if there would be no other mechanisms in the system, the host would still recognize this as an error and would have to wait for reinitialization by the (B)MC (the host is not allowed to write to the buffers in the transient state, i.e., until the data structures are reinitialized by (B)MC FW).
 3. Unexpected host reset (SW or HW reset is the same outcome): the host's unexpected reset will leave the data structures intact in (B)MC memory, so the (B)MC can still read the data from the buffers. Assuming the (B)MC understands the host's status via other mechanisms, the (B)MC can take informed decisions about how to respond to such situations.
- 488
- 489
- 490
- 491

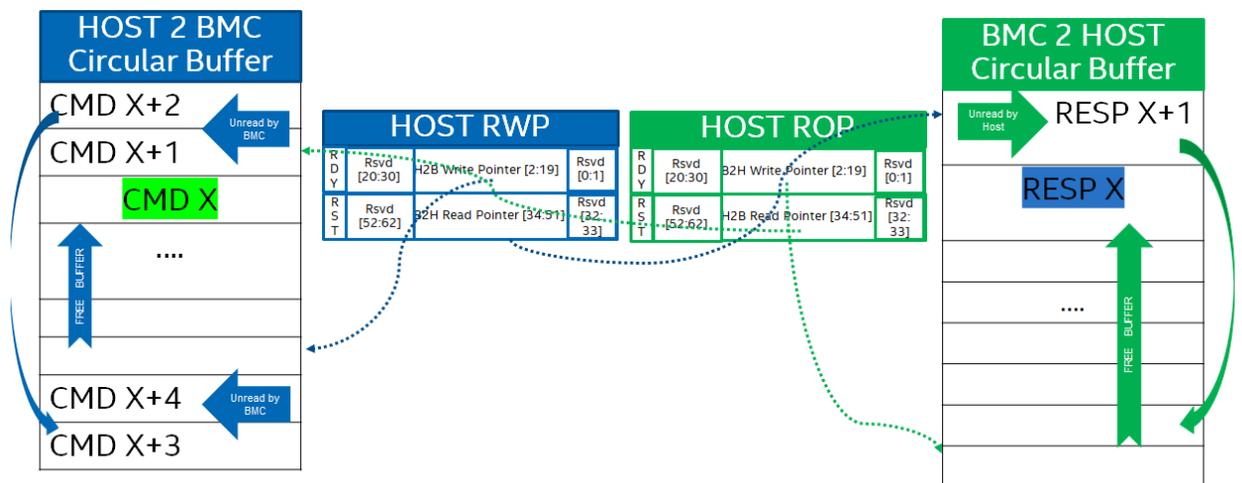
492 In all the above cases, MMBI data structures can be reinitialized after the reset to allow a clean restart.

493 **8.2 Calculation of Filled Space and Empty Space in Circular Buffer**

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

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

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



506

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

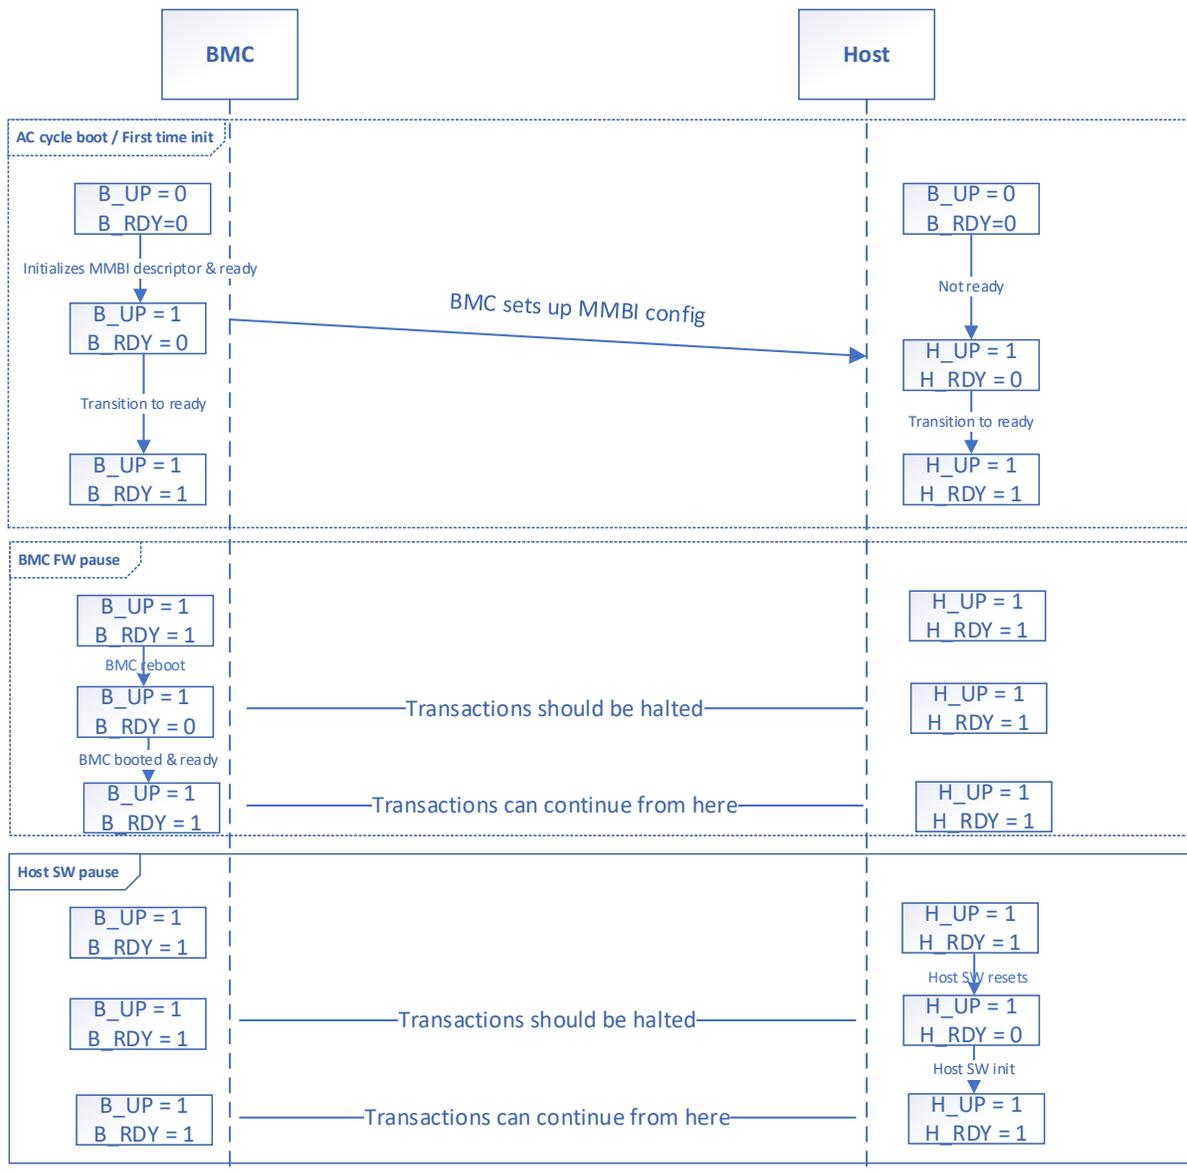
508 **8.3 Device Readiness and Communication Pause**

509 In addition to the reinitialization or reset states, the MMBI interface also uses the H_RDY and B_RDY
 510 flags to indicate the device’s readiness to consume incoming packets and handle them. When the host or
 511 (B)MC are ready to receive and handle packets, they set the B_RDY or H_RDY flags, respectively. If a
 512 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
 513 up but the target device is not ready to consume and handle new packets. When the interface is up, it
 514 means that the data structures are ready to accept new packets so the sender can:

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

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

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



524

525

Figure 9 – Sample MMBI Device Pause Sequences

526 **8.4 Packet Transfer**

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

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

564 8.5 Interrupts (Optional)

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

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

572 An interrupt handler shall:

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

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

581 9 Multi-Protocol Packet Format

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

586

Table 6 – Multi-Protocol Packet Format

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

587

588
589
590
591
592

ANNEX A (informative)

Notations

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

- 594 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets
595 starting from byte two and continuing to and including byte N. The lowest offset is on
596 the left; the highest is on the right.
- 597 • (6) Parentheses around a single number can be used in packet field descriptions to
598 indicate a byte field that may be present or absent.
- 599 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range
600 may be present or absent. The lowest offset is on the left; the highest is on the right.
- 601 • [PCle](#) Underlined blue text is typically used to indicate a reference to a document or
602 specification called out in clause 2, "Normative References" or to items hyperlinked
603 within the document.
- 604 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets
605 are given as zero-based values (that is, the least significant bit offset = 0).
- 606 • [7:5] A range of bit offsets. The most significant bit is on the left, the least significant bit is
607 on the right.
- 608 • 1b A number consisting of 0s and 1s followed by a lowercase "b" indicates that the
609 number is in binary format.
- 610 • 0x12A A leading "0x" indicates that the number is in hexadecimal format.

611
612
613
614
615

ANNEX B (informative)

Change log

Version	Date	Description
1.0.0	2023-07-14	Initial release
1.0.1	2024-08-30	Document title change ("Memory-Mapped BMC Interface" to "Memory-Mapped Buffer Interface") to better reflect potential broader uses of MMBI beyond just BMC
1.0.2	2025-07-11 2025-09-10	Define the discovery process of an MMBI description within a PCIe endpoint. Fixed: https://github.com/DMTF/PMCI-WG/issues/1646
1.1.0	2025-10-27 2025-10-28 2025-12-17	Enhanced the MMBI device with multi-channel capabilities to support concurrent applications and protocols. Fixed: https://github.com/DMTF/PMCI-WG/issues/1626 - Addressed the MMBI 1.1.0 RFC Ballot comments - Updated the RFC Ballot comments

616
617