



Document Identifier: DSP0234

Date: 2021-05-25

Version: 1.0.0

1
2
3
4

5 **CXL™ Fabric Manager API over MCTP Binding** 6 **Specification**

7 **Supersedes: None**
8 **Document Class: Normative**
9 **Document Status: Published**
10 **Document Language: en-US**

11

12 Copyright Notice

13 Copyright © 2021 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28 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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 patent may relate to or impact implementations of DMTF standards, visit
33 <http://www.dmtf.org/about/policies/disclosures.php>.

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

35

36

37

38

39

40

41

CONTENTS

42	Foreword	5
43	Introduction.....	6
44	Document conventions.....	6
45	1 Scope	7
46	2 Normative references	7
47	3 Terms and definitions.....	8
48	4 Symbols and abbreviated terms.....	8
49	5 Conventions	9
50	5.1 Reserved and unassigned values.....	9
51	5.2 Byte ordering.....	9
52	6 Overview.....	9
53	7 Message Type-specific considerations	10
54	7.1 Message Type number	10
55	7.2 CXL FM API over MCTP specification version information	10
56	7.3 Timing specifications.....	10
57	7.4 CXL FM API over MCTP message format.....	10
58	7.4.1 Integrity Check (IC), Tag Owner (TO) and Message Tag (Msg Tag) usage	11
59	7.4.2 Message assembly	11
60	7.5 Maximum message size	11
61	7.6 Multiple MCTP transports	12
62	ANNEX A (informative) Change log.....	13
63		

64 **Figures**

65 Figure 1: MCTP Message fields 10
66

67 **Tables**

68 Table 1: CXL FM API over MCTP Message field descriptions 11
69

70

Foreword

71 The CXL™ (*Compute Express Link™*) Fabric Manager API over MCTP Binding Specification (DSP0234)
72 was prepared by the Platform Management Components Intercommunications (PMCI Working Group) of
73 the DMTF.

74 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
75 management and interoperability. For information about the DMTF, see <http://www.dmtf.org>.

76 The CXL Consortium is an open industry, standards group formed to develop technical specifications that
77 facilitate breakthrough performance for emerging usage models while supporting an open ecosystem for
78 data center accelerators and other high-speed enhancements. For information about the CXL consortium,
79 see <https://www.computeexpresslink.org>.

80 Acknowledgments

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

82 Editors:

- 83 • Balaji Natrajan – Microchip Technology Inc
- 84 • Mahesh Natu – Intel Corporation

85 DMTF Contributors:

- 86 • Patrick Caporale – Lenovo
- 87 • Yuval Itkin – NVIDIA Corporation
- 88 • Eliel Louzoun – Intel Corporation
- 89 • Hemal Shah – Broadcom Inc.
- 90 • Bob Stevens – Dell Technologies

91 CXL Consortium Contributors:

- 92 • Vincent Hache – Microchip Technology Inc.
- 93 • Ariel Sibley – Microchip Technology Inc.

94

95

96

Introduction

97 The CXL™ Fabric Manager API over MCTP Binding Specification defines a new MCTP message type
98 used to convey CXL™ Fabric Manager API Messages over MCTP to devices.

99 Document conventions

100 Typographical conventions

101 This document uses the following typographical conventions:

- 102 • Document titles are marked in *italics*.
- 103 • Important terms that are used for the first time are marked in italics.
- 104 • Terms include a link to the term definition in the "Terms and definitions" clause, enabling easy
105 navigation to the term definition.
- 106 • ABNF rules are in monospaced font.

107 ABNF usage conventions

108 Format definitions in this document are specified using ABNF (see [RFC5234](#)), with the following
109 deviations:

- 110 • Literal strings are to be interpreted as case-sensitive Unicode characters, as opposed to the
111 definition in [RFC5234](#) that interprets literal strings as case-insensitive US-ASCII characters.
112

113 CXL™ Fabric Manager API over MCTP Binding Specification

114 1 Scope

115 The CXL™ Fabric Manager API over MCTP Binding Specification defines the bindings between CXL
116 Fabric Manager API protocol elements and MCTP elements in order to transport Fabric Manager API
117 Messages for CXL devices using MCTP. The specific Fabric Manager API message contents will be
118 documented outside of DMTF directly by the CXL consortium.

119 Portions of this specification rely on information and definitions from other specifications, which are
120 identified in clause 2. The following references are particularly relevant:

- 121 • DMTF DSP0236, *Management Component Transport Protocol (MCTP) Base Specification 1.3*,
122 defines the MCTP transport protocol over which the CXL Fabric Manager API over MCTP
123 messages are to be conveyed.
- 124 • CXL Consortium, *Compute Express Link™ (CXL™) Specification Revision 2.0*, defines the CXL
125 Fabric Manager API and message formats.

126 2 Normative references

127 The following referenced documents are indispensable for the application of this document. For dated or
128 versioned references, only the edition cited (including any corrigenda or DMTF update versions) applies.
129 For references without a date or version, the latest published edition of the referenced document
130 (including any corrigenda or DMTF update versions) applies. Earlier versions may not provide sufficient
131 support for this specification.

132 CXL Consortium, *Compute Express Link™ (CXL™) Specification Revision 2.0*,
133 <https://www.computeexpresslink.org>

134 DMTF DSP0236, *Management Component Transport Protocol (MCTP) Base Specification 1.3*
135 https://www.dmtf.org/sites/default/files/standards/documents/DSP0236_1.3.pdf

136 DMTF DSP0237, *Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding*
137 *Specification 1.2*
138 https://www.dmtf.org/sites/default/files/standards/documents/DSP0237_1.2.pdf

139 DMTF DSP0238, *Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding*
140 *Specification 1.1*
141 https://www.dmtf.org/sites/default/files/standards/documents/DSP0238_1.1.pdf

142 DMTF DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.7*
143 https://www.dmtf.org/sites/default/files/standards/documents/DSP0239_1.7.pdf

144 IETF, RFC4122, *A Universally Unique Identifier (UUID) URN Namespace*, July 2005
145 <http://www.ietf.org/rfc/rfc4122.txt>

146 IETF RFC5234, *ABNF: Augmented BNF for Syntax Specifications*, January 2008,
147 <http://tools.ietf.org/html/rfc5234>

148 DMTF DSP4004, *DMTF Release Process 2.4*,
149 http://dmtf.org/sites/default/files/standards/documents/DSP4004_2.4.pdf

150 **3 Terms and definitions**

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

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

159 The terms "clause", "subclause", "paragraph", and "annex" in this document are to be interpreted as
160 described in ISO/IEC Directives, Part 2, Clause 6.

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

164 Refer to [DSP0236](#) for terms and definitions that are used across the MCTP specifications.

165 Refer to the [CXL Specification](#) for terms and definitions that are used in the Compute Express Link™
166 Fabric Manager API specification. For the purposes of this document, the following additional terms and
167 definitions apply.

168 **3.1**

169 **Compute Express Link™**

170 A low-latency, high-bandwidth link that supports dynamic protocol muxing of coherent accesses, memory
171 access, and IO protocols, thus enabling attachment of coherent accelerators or memory devices.

172 **3.2**

173 **CXL™ Fabric Manager**

174 The Fabric Manager controls aspects of a CXL system related to binding and management of pooled
175 ports and devices.

176 **3.3**

177 **CXL™ Fabric Manager API**

178 Command set defined by the CXL consortium to manage devices in a CXL system.

179 **3.4**

180 **Endpoint**

181 An MCTP endpoint unless otherwise specified.

182 **4 Symbols and abbreviated terms**

183 Refer to [DSP0236](#) for terms and definitions that are used across the MCTP specifications. Refer to the
184 [CXL Specification](#) for terms and definitions that are used in the Compute Express Link™ Fabric Manager
185 API specification. For the purposes of this document, the following additional symbols and abbreviated
186 terms apply.

187 **4.1**

188 **CXL™**

189 Compute Express Link

190 **4.2**
191 **FM**
192 Fabric Manager

193 **4.3**
194 **MC**
195 Management Controller

196 **4.4**
197 **MCTP**
198 Management Component Transport Protocol

199 **5 Conventions**

200 **5.1 Reserved and unassigned values**

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

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

205 **5.2 Byte ordering**

206 Unless otherwise specified, the byte ordering of multibyte numeric fields or multibyte bit fields in this
207 specification shall be "Big Endian": The lowest byte offset holds the most significant byte and higher
208 offsets hold lesser significant bytes.

209 **6 Overview**

210 Compute Express Link™ (CXL) is a dynamic multi-protocol technology designed to support accelerators
211 and memory devices. CXL provides a rich set of protocols that include I/O semantics similar to PCIe (i.e.,
212 CXL.io), caching protocol semantics (i.e., CXL.cache), and memory access semantics (i.e., CXL.mem)
213 over a discrete or on-package link.

214 CXL devices can be configured statically or dynamically via a Fabric Manager (FM), an external logical
215 process that queries and configures the system's operational state by using the Fabric Manager
216 Application Programming Interface (FM API) commands. FM API commands are defined by the members
217 of CXL Consortium in the Compute Express Link specification. Refer to www.computeexpresslink.org and
218 the [CXL specification](#) for more information.

219 This specification only defines how FM API Commands are encapsulated in MCTP Messages and
220 transferred between MCTP Endpoints over transports that have a corresponding MCTP transport binding
221 specification. These are referred to in this document as FM API Messages over MCTP. The definitions
222 and semantics of the FM API Commands themselves are outside the scope of this specification and are
223 defined in the CXL specification.

224 The MCTP Transport Bindings that are used for CXL FM API over MCTP are defined in other companion
225 specifications including but not limited to the MCTP SMBus/I2C Transport Binding Specification
226 ([DSP0237](#)) and the MCTP PCIe VDM Transport Binding Specification ([DSP0238](#)).

227 **7 Message Type-specific considerations**

228 **7.1 Message Type number**

229 The Message Type number for CXL FM API over MCTP messages is defined in the MCTP IDs and
 230 Codes Specification ([DSP0239](#)) and the number assigned is 0x07.

231 **7.2 CXL FM API over MCTP specification version information**

232 Implementations that follow this specification shall return the following version information in the response
 233 to the GET MCTP Version Support command when the Message Type parameter in the request is set to
 234 0x07 (return CXL FM API over MCTP specification version information).

235 The Version Number Entry 1 field shall be used to indicate compatibility with Version 1.0.0 of the CXL FM
 236 API over MCTP message type as:

237 1.0 [Major version 1, minor version 0, any update version, no alpha]

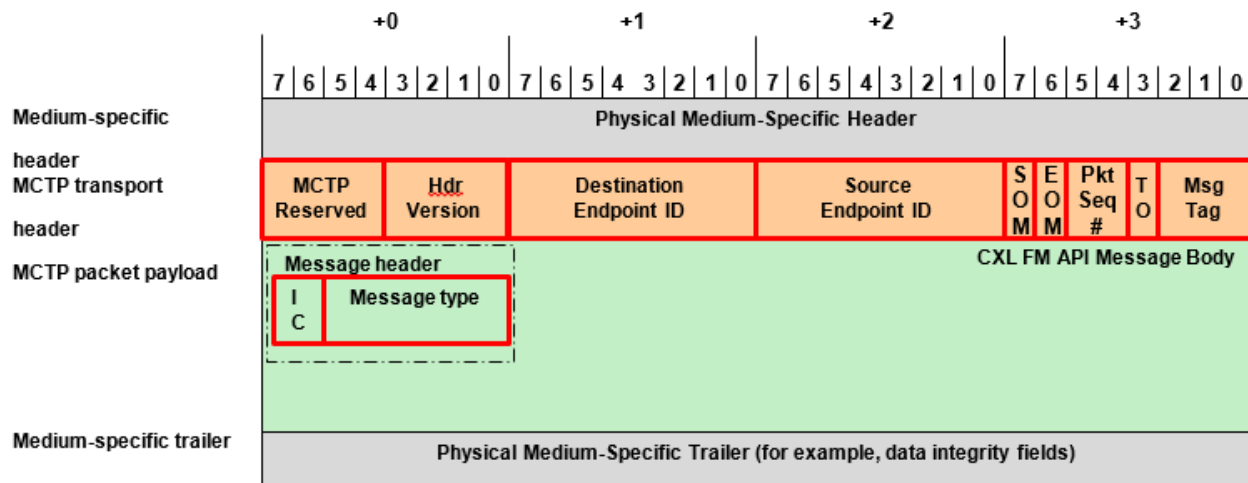
238 This is reported using the encoding as: 0xF1F0FF00

239 **7.3 Timing specifications**

240 CXL FM API messages over MCTP are made up of one or more MCTP packets. Each MCTP packet shall
 241 comply with the timing, arbitration, and fairness requirements of the transport binding specifications for
 242 the media through which it passes. For examples, refer to the MCTP SMBus/I2C Transport Binding
 243 Specification ([DSP0237](#)) and the MCTP PCIe VDM Transport Binding Specification ([DSP0238](#)) for
 244 specific packet and message timing requirements.

245 **7.4 CXL FM API over MCTP message format**

246 Referring to Figure 1, the CXL FM API Messages over MCTP are carried via the MCTP packet payload of
 247 one or more MCTP packets.



248 **Figure 1: MCTP Message fields**

250 **7.4.1 Integrity Check (IC), Tag Owner (TO) and Message Tag (Msg Tag) usage**

251 CXL FM API over MCTP messages do not include an overall message integrity check field and shall set
 252 the Integrity Check (IC) bit to 0b.

253 CXL FM API request and event notification messages when transported over MCTP shall have the TO bit
 254 set (TO bit = 1b). CXL FM API response messages over MCTP shall have the TO bit cleared (TO bit =
 255 0b).

256 When request/response message exchange is used and the Tag Owner (TO) bit is set to 1b in the
 257 request, a responder shall return the same Message Tag with the Tag Owner bit cleared to 0b in the
 258 corresponding response message.

259

Field Name	Field Size	Description
TO	1 bit	1b - CXL FM API Request and Event Notification messages 0b – CXL FM API Response messages
IC	1 bit	Message Integrity Check bit = 0b CXL FM API over MCTP messages do not include an overall Message Integrity check field.
Message type	7 bits	CXL FM API = 0x07 This field identifies the MCTP message as carrying a CXL FM API message.
CXL FM API Message Body	Variable	The CXL FM API message fields are defined in the CXL Specification

260 **Table 1: CXL FM API over MCTP Message field descriptions**

261 For a definition of CXL FM API request, response and event notification messages, refer to the [CXL](#)
 262 [Specification](#).

263 **7.4.2 Message assembly**

264 CXL FM API messages over MCTP may be split into one or more MCTP packets thus requiring
 265 segmentation and assembly. All multi-packet CXL FM API over MCTP messages shall comply with the
 266 message packetization and assembly rules of the MCTP base specification ([DSP236](#)). Specifically,
 267 clauses in the MCTP base specification related to Message assembly, Dropped packets, Starting
 268 message assembly, Terminating message assembly/dropped messages, and Dropped messages shall
 269 be complied with strictly. CXL FM API messages when transported over MCTP shall not require any
 270 changes to the MCTP base specification.

271 **7.5 Maximum message size**

272 The CXL FM API message body over MCTP shall be less than or equal to 1088 bytes. All MCTP endpoint
 273 shall support this maximum message body size of 1088 bytes, which includes a maximum of 1024 bytes
 274 for the CXL FM API message payload and a maximum of 64 bytes for the CXL FM API header. This
 275 corresponds to a transfer of 17 MCTP packets using a baseline transmission unit of 64 bytes for the
 276 MCTP packet payload. See the CXL Specification for a definition of the CXL FM API message payload
 277 and headers.

278 The maximum message size includes the IC bit and Message Type fields plus any additional Message
 279 Type-specific header fields, as required by the CXL FM API. Refer to the [CXL Specification](#) for any
 280 additional restrictions on message sizes.

281 7.6 Multiple MCTP physical transports

282 In order to facilitate identification of devices that are accessible via multiple physical transports, the
283 endpoints in the device shall support the Get Endpoint UUID MCTP command.

284 An MCTP endpoint is not required to support more than one outstanding command over a single physical
285 transport. A requestor shall not have multiple requests outstanding simultaneously across multiple
286 physical transports to an endpoint. Otherwise, this specification does not define any additional behaviors
287 related to communicating with CXL™ devices over MCTP that may be accessed through more than one
288 type of MCTP physical transport on a given MCTP network.

289
290
291
292

ANNEX A (informative)

Change log

Version	Date	Description
1.0.0	2021-05-25	

293