



Document Identifier: DSP0236

Date: 2026-02-13

Version: 1.4.0WIP80

Management Component Transport Protocol (MCTP) Base Specification

Includes MCTP Control Specifications

Information for Work-in-Progress version:

IMPORTANT: This document is not a standard. It does not necessarily reflect the views of DMTF or its members. Because this document is a Work in Progress, this document may still change, perhaps profoundly and without notice. This document is available for public review and comment until superseded.

Provide any comments through the DMTF Feedback Portal:

<https://www.dmtf.org/standards/feedback>

Supersedes: 1.3.3

Document Class: Normative

Document Status: Work in Progress

Document Language: en-US

12 Copyright notice

13 Copyright © 2026 Distributed Management Task Force, Inc. (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,
22 or identify any or all such third-party patent right, owners or claimants, nor for any incomplete or
23 inaccurate identification or disclosure of such rights, owners or claimants. DMTF shall have no liability to
24 any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize,
25 disclose, or identify any such third-party patent rights, or for such party's reliance on the standard or
26 incorporation thereof in its product, protocols or testing procedures. DMTF shall have no liability to any
27 party implementing such standard, 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29 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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 the DMTF that, in their opinion,
33 such patent may relate to or impact implementation of the DMTF standards, visit
34 <http://www.dmtf.org/about/policies/disclosures.php>

35 PCI-SIG, PCIe, and the PCI HOT PLUG design mark are registered trademarks or service marks of PCI-
36 SIG.

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

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

CONTENTS

40 Foreword 7

41 Introduction 8

42 1 Scope 9

43 2 Normative references 9

44 2.1 Approved references 9

45 2.2 Other references 9

46 3 Terms and definitions 10

47 3.1 Requirement term definitions 10

48 3.2 MCTP term definitions 12

49 4 Symbols and abbreviated terms 18

50 5 Conventions 20

51 5.1 Overview 20

52 5.2 Byte ordering 20

53 5.3 Reserved fields 20

54 6 Management component relationships 21

55 7 MCTP overview 21

56 8 MCTP base protocol 24

57 8.1 Overview 24

58 8.2 MCTP packet fields 24

59 8.3 Special endpoint IDs 26

60 8.4 Packet payload and transmission unit sizes 27

61 8.5 Maximum message body sizes 27

62 8.6 Message assembly 28

63 8.7 Dropped packets 28

64 8.8 Starting message assembly 29

65 8.9 Terminating message assembly/dropped messages 29

66 8.10 Dropped messages 30

67 8.11 MCTP versioning and message type support 30

68 8.12 MCTP message types 32

69 8.13 Security 32

70 8.14 Limitations 32

71 8.15 MCTP discovery and addressing 33

72 8.16 Devices with multiple media interfaces 34

73 8.17 Peer transactions 34

74 8.18 Endpoint ID assignment and endpoint ID pools 35

75 8.19 Handling reassigned EIDs 40

76 9 MCTP bridging 41

77 9.2 Bridge and routing table examples 49

78 9.3 Endpoint ID resolution 53

79 9.4 Bridge and bus owner implementation recommendations 55

80 9.5 Path and transmission unit discovery 56

81 9.6 Path transmission unit requirements for bridges 60

82 10 Rate Limiting 60

83 11 Resiliency and recovery 63

84 11.1 Overview 63

85 11.2 MCTP reset types 64

86 11.3 MCTP System Resiliency 64

87 12 MCTP control protocol 65

88 12.2 Terminology 65

89 12.3 Control message classes 66

90 12.4 MCTP control message format 66

91	12.5	MCTP control message fields.....	67
92	12.6	MCTP control message transmission unit size.....	68
93	12.7	Tag Owner (TO), Request (Rq), and Datagram (D) bit usage.....	68
94	12.8	Concurrent command processing.....	69
95	13	MCTP control messages.....	69
96	13.1	Overview.....	69
97	13.2	MCTP control message command codes.....	70
98	13.3	MCTP control message completion codes.....	72
99	13.4	Set Endpoint ID.....	73
100	13.5	Get Endpoint ID.....	75
101	13.6	Get Endpoint UUID.....	76
102	13.7	Get MCTP version support.....	77
103	13.8	Get Message Type Support.....	80
104	13.9	Get Vendor Defined Message Support.....	80
105	13.10	Resolve Endpoint ID.....	82
106	13.11	Allocate Endpoint IDs.....	83
107	13.12	Routing Information Update.....	85
108	13.13	Get Routing Table Entries.....	87
109	13.14	Prepare for Endpoint Discovery.....	88
110	13.15	Endpoint Discovery.....	89
111	13.16	Discovery Notify.....	89
112	13.17	Get Network ID.....	90
113	13.18	Query Hop.....	90
114	13.19	Resolve UUID.....	91
115	13.20	Query rate limit.....	92
116	13.21	Request TX rate limit.....	93
117	13.22	Update rate limit.....	94
118	13.23	Query Supported Interfaces.....	94
119	13.24	Query Endpoint Advanced Capabilities.....	95
120	13.25	Request EID pool.....	95
121	13.26	Release EID pool.....	96
122	13.27	Query Supported Recovery Actions.....	96
123	13.28	Recovery Request.....	97
124	13.29	Transport Specific.....	98
125	14	Vendor Defined – PCI and Vendor Defined – IANA messages.....	99
126	14.1	Vendor Defined – PCI message format.....	100
127	14.2	Vendor Defined – IANA message format.....	100
128		ANNEX A (informative) Notation.....	101
129		ANNEX B (informative) Change log.....	102
130			

131 Figures

132	Figure 1 – Management component relationships.....	21
133	Figure 2 – MCTP networks.....	22
134	Figure 3 – MCTP topology.....	23
135	Figure 4 – Generic message fields.....	24
136	Figure 5 – Topmost bus owners.....	36
137	Figure 6 – Split bridge.....	36
138	Figure 7 – Acceptable failover/redundant communication topologies.....	42
139	Figure 8 – Routing/bridging restrictions.....	42
140	Figure 9 – EID options for MCTP bridges.....	43

141 Figure 10 – Basic routing table entry fields 46
 142 Figure 11 – Routing table population 47
 143 Figure 12 – Example 1 Routing topology 49
 144 Figure 13 – Example 2 Routing topology 51
 145 Figure 14 – Example 3 Routing topology 52
 146 Figure 15 – Endpoint ID resolution 54
 147 Figure 16 – Resolving multiple paths 55
 148 Figure 17 – Example path routing topology 57
 149 Figure 18 – Path transmission unit discovery flowchart 60
 150 Figure 19 – Example rate limiting message exchanges 61
 151 Figure 21 – MCTP control message format 67
 152 Figure 22 – Structure of Vendor ID field for Get Vendor Defined capabilities message 81
 153 Figure 23 – EID Pools from multiple bus owners 84

154 **Tables**

155 Table 1 – MCTP base protocol common fields 24
 156 Table 2 – Special endpoint IDs 26
 157 Table 3 – MCTP Message Types Used in this Specification 32
 158 Table 4 – Example 1 Routing table for D2 50
 159 Table 5 – Example 2 Routing table for D1 51
 160 Table 6 – Example 3 Routing table for D2 52
 161 Table 7 – Additional information tracked by bridges 53
 162 Table 8 – MCTP control protocol terminology 65
 163 Table 9 – MCTP control message types 66
 164 Table 10 – MCTP control message fields 67
 165 Table 11 – Tag Owner (TO), Request (Rq) and Datagram (D) bit usage 68
 166 Table 12 – MCTP control command numbers 70
 167 Table 13 – MCTP control message completion codes 72
 168 Table 14 – Set Endpoint ID message 73
 169 Table 15 – Get Endpoint ID message 75
 170 Table 16 – Get Endpoint UUID message format 76
 171 Table 17 – Example UUID format 76
 172 Table 18 – Get MCTP version support message 77
 173 Table 19 – Get Message Type Support message 80
 174 Table 20 – Get Vendor Defined Message Support message 81
 175 Table 21 – Vendor ID formats 82
 176 Table 22 – Resolve Endpoint ID message 82
 177 Table 23 – Allocate Endpoint IDs message 84
 178 Table 24 – Routing Information Update message 86
 179 Table 25 – Routing Information Update entry format 86
 180 Table 26 – Get Routing Table Entries message 87
 181 Table 27 – Routing Table Entry format 87
 182 Table 28 – Prepare for Endpoint Discovery message 89
 183 Table 29 – Endpoint Discovery message 89
 184 Table 30 – Discovery Notify message 90
 185 Table 31 – Get Network ID message format 90

186	Table 32 – Query Hop message	90
187	Table 33 – Resolve UUID message.....	92
188	Table 34 – Resolve UUID message entry format	92
189	Table 35 – Query rate limit message	92
190	Table 36 – Request TX rate limit message.....	93
191	Table 37 – Update rate limit message	94
192	Table 38 – Query Supported Interfaces.....	95
193	Table 39 – Query Endpoint Advanced Capabilities	95
194	Table 40 – Request EID pool.....	95
195	Table 41 – EID section structure.....	96
196	Table 42 – Release EID pool	96
197	Table 43 – Query Supported Recovery Actions	97
198	Table 44 Reset levels.....	97
199	Table 45 –Recovery Request	98
200	Table 46 – Transport Specific message	98
201	Table 47 – Vendor Defined – PCI message format	100
202	Table 48 – Vendor Defined – IANA message format.....	100
203		

204

Foreword

205 The *Management Component Transport Protocol (MCTP) Base Specification* (DSP0236) was prepared
206 by the PMCI Working Group.

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

209 This version supersedes version 1.3.3. For a list of changes, see the change log in ANNEX B.

210 Acknowledgments

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

212 Editor:

213 Yuval Itkin – NVIDIA Corporation

214 Contributors:

- 215 • Alan Berenbaum – SMSC
- 216 • Patrick Caporale – Lenovo
- 217 • Phil Chidester – Dell Inc
- 218 • Ira Kalman – Intel Corporation
- 219 • Edward Klodnicki - IBM
- 220 • Joe Kozlowski – Dell Inc
- 221 • Patrick Kutch – Intel Corporation
- 222 • John Leung - Intel Corporation
- 223 • Eliel Louzoun – Intel Corporation
- 224 • Patrick Schoeller – Hewlett Packard Enterprise
- 225 • Hemal Shah - Broadcom Limited
- 226 • Tom Slaight – Intel Corporation
- 227 • Bob Stevens – Dell Inc.

228

Introduction

229 The Management Component Transport Protocol (MCTP) defines a communication model intended to
230 facilitate communication between:

- 231 • Management controllers and other management controllers
- 232 • Management controllers and managed devices

233 The communication model includes a message format, transport description, message exchange
234 patterns, and configuration and initialization messages.

235 MCTP is designed so that it can potentially be used on many bus types. The protocol is intended to be
236 used for intercommunication between elements of platform management subsystems used in computer
237 systems, and is suitable for use in mobile, desktop, workstation, and server platforms. Management
238 controllers such as a baseboard management controller (BMC) can use this protocol for communication
239 between one another, as well as for accessing managed devices within the platform.

240 Management controllers can use this protocol to send and receive MCTP-formatted messages across the
241 different bus types that are used to access managed devices and other management controllers.
242 Managed devices in a system need to provide an implementation of the message format to facilitate
243 actions performed by management controllers.

244 It is intended that different types of devices in a management system might need to implement different
245 portions of the complete capabilities defined by this protocol. Where relevant, this is called out in the
246 individual requirements.
247

Management Component Transport Protocol (MCTP) Base Specification

1 Scope

The *MCTP Base Specification* describes the command protocol, requirements, and use cases of a transport protocol for communication between discrete management controllers on a platform, as well as between management controllers and the devices they manage.

This document is intended to meet the following objectives:

- Describe the MCTP Base transport protocol.
- Describe the MCTP control message protocol.

The MCTP specifies a transport protocol format. This protocol is independent of the underlying physical bus properties, as well as the "data-link" layer messaging used on the bus. The physical and data-link layer methods for MCTP communication across a given medium are defined by companion "transport binding" specifications, such as [DSP0238](#), MCTP over PCIe® Vendor Defined Messaging, and [DSP0237](#), MCTP over SMBus/I²C. This approach enables future transport bindings to be defined to support additional buses such as USB, RMII, and others, without affecting the base MCTP specification.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

2.1 Approved references

DMTF DSP4004, *DMTF Release Process v2.7*

http://www.dmtf.org/standards/published_documents/DSP4004_2.7.pdf

DMTF DSP2016, Management Component Transport Protocol (MCTP) Overview White Paper

http://www.dmtf.org/standards/published_documents/DSP2016_1.0.pdf

DMTF, DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes*

http://www.dmtf.org/standards/published_documents/DSP0239_1.3.x.pdf

DMTF DSP0237, Management Component Transport Protocol SMBus/I²C Transport Binding Specification

http://www.dmtf.org/standards/published_documents/DSP0237_1.0.x.pdf

DMTF DSP0238, Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding Specification

http://www.dmtf.org/standards/published_documents/DSP0238_1.0.x.pdf

2.2 Other references

Hewlett-Packard, Intel, Microsoft, Phoenix, and Toshiba, *Advanced Configuration and Power Interface Specification v5.0*, ACPI, December 6, 2011

<http://www.acpi.info/downloads/ACPIspec50.pdf>

- 284 IETF, RFC20, *ASCII format for Network Interchange*, October 16, 1969
285 <http://tools.ietf.org/html/rfc20>
- 286 IETF, RFC4122, *A Universally Unique Identifier (UUID) URN Namespace*, July 2005
287 <http://datatracker.ietf.org/doc/rfc4122/>
- 288 IETF, RFC2119, *Key Words for use in RFCs to Indicate Requirement Levels*, March 1997
289 <http://datatracker.ietf.org/doc/rfc2119/>
- 290 Intel, Hewlett-Packard, NEC, and Dell, *Intelligent Platform Management Interface Specification: Second
291 Generation v2.0*, IPMI, 2004
292 <http://www.intel.com/design/servers/ipmi>
- 293 ISO/IEC Directives, Part 2, *Principles and Rules for the structure and drafting of ISO and IEC documents*
294 <https://www.iso.org/sites/directives/current/part2/index.xhtml>
- 295 PCI-SIG, PCI Express™ Specifications
296 <http://www.pcisig.com/specifications/pciexpress/>
- 297 NXP Semiconductors, *UM10204 I2C-bus specification and user manual*, Rev. 5, October 9, 2012
298 http://www.nxp.com/documents/user_manual/UM10204.pdf
- 299 SMBus, *System Management Bus (SMBus) Specification v2.0*, SMBus, 2000
300 <http://www.smbus.org/specs/smbus20.pdf>

301 **3 Terms and definitions**

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

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

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

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

315 **3.1 Requirement term definitions**

316 This clause defines key phrases and words that denote requirement levels in this specification. These
317 definitions are consistent with the terms defined in [RFC2119](#).

318 **3.1.1**

319 **can**

320 used for statements of possibility and capability, whether material, physical, or causal

321 **3.1.2**

322 **cannot**

323 used for statements of possibility and capability, whether material, physical or causal

- 324 **3.1.3**
325 **conditional**
326 indicates requirements to be followed strictly to conform to the document when the specified conditions
327 are met.
- 328 **3.1.4**
329 **deprecated**
330 indicates that an element or profile behavior has been outdated by newer constructs.
- 331 **3.1.5**
332 **mandatory**
333 indicates requirements to be followed strictly to conform to the document and from which no deviation is
334 permitted.
- 335 **3.1.6**
336 **may**
337 indicates a course of action permissible within the limits of the document.
- 338 Note 1 to entry: An implementation that does *not* include a particular option shall be prepared to interoperate with
339 another implementation that *does* include the option, although perhaps with reduced functionality. An implementation
340 that *does* include a particular option shall be prepared to interoperate with another implementation that does *not*
341 include the option (except for the feature that the option provides).
- 342 **3.1.7**
343 **may not**
344 indicates flexibility of choice with no implied preference
- 345 **3.1.8**
346 **need not**
347 indicates a course of action permissible within the limits of the document.
- 348 **3.1.9**
349 **not recommended**
350 indicates that valid reasons may exist in particular circumstances when the particular behavior is
351 acceptable or even useful, but the full implications should be understood and carefully weighed before
352 implementing any behavior described with this label.
- 353 **3.1.10**
354 **obsolete**
355 indicates that an item was defined in prior specifications but has been removed from this specification.
- 356 **3.1.11**
357 **optional**
358 indicates a course of action permissible within the limits of the document.
- 359 **3.1.12**
360 **physical layer**
361 the electrical bus of the MCTP network.

362 **3.1.13**363 **recommended**

364 indicates that valid reasons may exist in particular circumstances to ignore a particular item, but the full
365 implications should be understood and carefully weighed before choosing a different course.

366 **3.1.14**367 **required**

368 indicates that the item is an absolute requirement of the specification.

369 **3.1.15**370 **shall**

371 indicates that the item is an absolute requirement of the specification.

372 **3.1.16**373 **shall not**

374 indicates that the definition is an absolute prohibition of the specification.

375 **3.1.17**376 **should**

377 indicates that among several possibilities, one is recommended as particularly suitable, without
378 mentioning or excluding others, or that a certain course of action is preferred but not necessarily required.

379 **3.1.18**380 **should not**

381 indicates that a certain possibility or course of action is deprecated but not prohibited.

382 **3.2 MCTP term definitions**

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

384 **3.2.1**385 **Address Resolution Protocol**386 **ARP**

387 refers to the procedure used to dynamically determine the addresses of devices on a shared
388 communication medium.

389 **3.2.2**390 **baseline transmission unit**391 **BTU**

392 the required common denominator size of a transmission unit for packet payloads that are carried in an
393 MCTP packet. Baseline Transmission Unit-sized packets are guaranteed to be routable within an MCTP
394 network.

395 **3.2.3**396 **baseboard management controller**397 **BMC**

398 **3.2.4** a term coined by the IPMI specifications for the main management controller in an IPMI-based
399 platform management subsystem. Also sometimes used as a generic name for a motherboard resident

400 management controller that provides motherboard-specific hardware monitoring and control functions for
401 the platform management subsystem.

402 **binary-coded decimal**

403 **BCD**

404 indicates a particular binary encoding for decimal numbers where each four bits (*nibble*) in a binary
405 number is used to represent a single decimal digit, and with the least significant four bits of the binary
406 number corresponding to the least significant decimal digit. The binary values 0000b through 1001b
407 represent decimal values 0 through 9, respectively.

408 EXAMPLE: BCD encoding a byte can represent a two-digit decimal number where the most significant
409 nibble (bits 7:4) of the byte contains the encoding for the most significant decimal digit and the least
410 significant nibble (bits 3:0) contains the encoding for the least significant decimal digit (for example,
411 0010_1001b in BCD encoding corresponds to the decimal number 29).

412 **3.2.5**

413 **bridge**

414 generically, the circuitry and logic that connects one computer bus or interconnect to another, allowing an
415 agent on one to access the other. Within this document, the term *bridge* refers to MCTP bridge, unless
416 otherwise indicated.

417 **3.2.6**

418 **burst**

419 a number of consecutive baseline transmission unit Packets that the transmitter endpoint sends with
420 minimal delay between those baseline transmission unit packets.

421 **3.2.7**

422 **bus**

423 a physical addressing domain shared between one or more platform components that share a common
424 physical layer address space.

425 **3.2.8**

426 **bus owner**

427 the party responsible for managing address assignments (can be logical or physical addresses) on a bus
428 (for example, in MCTP, the bus owner is the party responsible for managing EID assignments for a given
429 bus). A bus owner may also have additional media-specific responsibilities, such as assignment of
430 physical addresses.

431 **3.2.9**

432 **byte**

433 an 8-bit quantity. Also referred to as an *octet*.

434 Note 1 to entry: PMCI specifications use the term *byte*, not *octet*.

435 **3.2.10**

436 **endpoint**

437 see [MCTP endpoint](#)

438 **3.2.11**

439 **endpoint ID**

440 **EID**

441 see [MCTP endpoint ID](#)

442 **3.2.12**443 **Globally Unique Identifier**444 **GUID**445 see [UUID](#)446 **3.2.13**447 **host interface**

448 a hardware interface and associated protocols that is used by software running locally on the host
449 processors to access the hardware of a management subsystem within a managed system.

450 **3.2.14**451 **Inter-Integrated Circuit**452 **I²C**

453 a multi-master, two-wire, serial bus originally developed by Philips Semiconductor; now maintained by
454 NXP Semiconductors,

455 **3.2.15**456 **intelligent management device**457 **IMD**

458 a management device that is typically implemented using a microcontroller and accessed through a
459 messaging protocol. Management parameter access provided by an IMD is typically accomplished using
460 an abstracted interface and data model rather than through direct "register level" accesses.

461 **3.2.16**462 **Intelligent Platform Management Interface**463 **IPMI**

464 a set of specifications defining interfaces and protocols originally developed for server platform
465 management by the IPMI Promoters Group: Intel, Dell, HP, and NEC

466 **3.2.17**467 **managed entity**

468 the physical or logical entity that is being managed through management parameters.

469 EXAMPLE 1: *physical* entities include fans, processors, power supplies, circuit cards, chassis, and so on.

470 EXAMPLE 2: *logical* entities include virtual processors, cooling domains, system security states, and so
471 on.

472 **3.2.18**473 **Management Component Transport Protocol**474 **MCTP**

475 The protocol defined in this specification.

476 **3.2.19**477 **management controller**

478 a microcontroller or processor that aggregates management parameters from one or more managed
479 devices and makes access to those parameters available to local or remote software, or to other
480 management controllers, through one or more management data models. Management controllers may
481 also interpret and process management-related data, and initiate management-related actions on
482 managed devices. While a native data model is defined for PMCI, it is designed to be capable of
483 supporting other data models, such as CIM, IPMI, and vendor-specific data models. The microcontroller
484 or processor that serves as a management controller can also incorporate the functions of a management
485 device.

486 **3.2.20**487 **managed device**

488 for this specification, managed device refers to a device that is typically implemented using a
489 microcontroller and accessed through a messaging protocol and is used for accessing one or more
490 management parameters. Management parameter access provided by a managed device is typically
491 accomplished using an abstracted interface and data model rather than through direct "register level"
492 accesses. A managed device responds to management requests, but does not initiate or aggregate
493 management operations except in conjunction with a management controller (that is, it is a *satellite*
494 device that is subsidiary to one or more management controllers).

495 **3.2.21**496 **management parameter**

497 a particular datum representing a characteristic, capability, status, or control point associated with a
498 managed entity. Example management parameters include temperature, speed, volts, on/off, link state,
499 uncorrectable error count, device power state, and so on.

500 **3.2.22**501 **MCTP bridge**

502 an MCTP endpoint that can route MCTP messages not destined for itself that it receives on one
503 interconnect onto another without interpreting them. The ingress and egress media at the bridge may be
504 either homogeneous or heterogeneous. Also referred to in this document as a "bridge".

505 **3.2.23**506 **MCTP bus owner**

507 responsible for EID assignment for MCTP or translation on the buses that it is a master of. The MCTP bus
508 owner may also be responsible for physical address assignment. For example, for SMBus/I2C bus
509 segments, the MCTP bus owner is also the ARP master. This means the bus owner assigns dynamic
510 SMBus/I2C addresses to those devices requiring it.

511 **3.2.24**512 **MCTP control command**

513 commands defined under the MCTP *control* message type that are used for the initialization and
514 management of MCTP communications (for example, commands to assign EIDs, discover device MCTP
515 capabilities, and so on)

516 **3.2.25**517 **MCTP endpoint**

518 an MCTP communication terminus

519 Note 1 to entry: An MCTP endpoint is a terminus or origin of MCTP packets or messages. That is, the
520 combined functionality within a physical device that communicates using the MCTP transport protocol and
521 handles MCTP control commands. This includes MCTP-capable management controllers and managed
522 devices.

523 **3.2.26 Note 2 to entry: Also referred to in this document as "endpoint".**524 **MCTP endpoint ID**

525 the logical address used to route MCTP messages to a specific MCTP endpoint.

526 Note 1 to entry: A numeric handle (logical address) that uniquely identifies a particular MCTP endpoint
527 within a system for MCTP communication and message routing purposes. Endpoint IDs are unique
528 among MCTP endpoints that comprise an MCTP communication network within a system. MCTP EIDs
529 are only unique within a particular MCTP network. That is, they can be duplicated or overlap from one
530 MCTP network to the next.

531 **3.2.27 Note 2 to entry: Also referred to in this document as "endpoint ID" and abbreviated "EID".**
532 **MCTP host interface**

533 a host interface that enables host software to locally access an MCTP Network in the managed system.

534 **3.2.28**

535 **MCTP management controller**

536 a management controller that is an MCTP endpoint

537 Note 1 to entry: Unless otherwise indicated, the term "management controller" refers to an "MCTP
538 management controller" in this document.

539 **3.2.29**

540 **MCTP managed device**

541 a managed device that is an MCTP endpoint

542 Note 1 to entry: Unless otherwise indicated, the term "managed device" refers to an "MCTP managed
543 device" in this document.

544 **3.2.30**

545 **MCTP message**

546 a unit of communication based on the message type that is relayed through the MCTP Network using one
547 or more MCTP packets.

548 **3.2.31**

549 **MCTP network**

550 a collection of MCTP endpoints that communicate using MCTP and share a common MCTP endpoint ID
551 space.

552 **3.2.32**

553 **MCTP network ID**

554 a unique identifier to distinguish each independent MCTP network within a platform.

555 **3.2.33**

556 **MCTP packet**

557 the unit of data transfer used for MCTP communication on a given physical medium.

558 **3.2.34**

559 **MCTP packet payload**

560 refers to the portion of the message body of an MCTP message that is carried in a single MCTP packet.

561 **3.2.35**

562 **message**

563 see [MCTP message](#)

564 **3.2.36**

565 **message assembly**

566 the process of receiving and linking together two or more MCTP packets that belong to a given MCTP
567 message to allow the entire message header and message data (payload) to be extracted.

568 **3.2.37**

569 **message body**

570 the portion of an MCTP message that carries the message type field and any message type-specific data
571 associated with the message.

572 Note 1 to entry: An MCTP message spans multiple MCTP packets when the message body needs are
573 larger than what can fit in a single MCTP packet. Thus, the message body portion of an MCTP message
574 can span multiple MCTP packets.

575 **3.2.38**
576 **message disassembly**

577 the process of taking an MCTP message where the message's header and data (payload) cannot be
578 carried in a single MCTP packet and generating the sequence of two or more packets required to deliver
579 that message content within the MCTP network.

580 **3.2.39**
581 **message originator**

582 the original transmitter (source) of a message targeted to a particular message terminus.

583 **3.2.40**
584 **message terminus**

585 the name for a triplet of fields called the MCTP Source Endpoint ID, Tag Owner bit value, and Message
586 Tag value.

587 Note 1 to entry: Together, these fields identify the packets for an MCTP message within an MCTP
588 network for the purpose of message assembly. The message terminus itself can be thought of as
589 identifying a set of resources within the recipient endpoint that is handling the assembly of a particular
590 message.

591 **3.2.41**
592 **most significant byte**
593 **MSB**

594 refers to the highest order byte in a number consisting of multiple bytes.

595 **3.2.42**
596 **nibble**

597 the computer term for a four-bit aggregation, or half of a byte

598 **3.2.43**
599 **packet**

600 see [MCTP packet](#)

601 **3.2.44**
602 **packet payload**

603 see [MCTP packet payload](#)

604 **3.2.45**
605 **payload**

606 refers to the information bearing fields of a message.

607 Note 1 to entry: This is separate from those fields and elements that are used to transport the message
608 from one point to another, such as address fields, framing bits, checksums, and so on. In some instances,
609 a given field may be both a payload field and a transport field.

610 **3.2.46**
611 **physical transport binding**

612 refers to specifications that define how the MCTP base protocol and MCTP control commands are
613 implemented on a particular physical transport type and medium, such as SMBus/I²C, PCI Express™
614 Vendor Defined Messaging, and so on.

615 **3.2.47**

616 **Platform Management Component Intercommunications**

617 **PMCI**

618 name for a working group under the Distributed Management Task Force's Pre-OS Workgroup that is
619 chartered to define standardized communication protocols, low level data models, and transport
620 definitions that support communications with and between management controllers and managed devices
621 that form a platform management subsystem within a managed computer system.

622 **3.2.48**

623 **Rate Limiting**

624 a method for limiting the data rate sent from an MCTP endpoint to another MCTP endpoint.

625 **3.2.49**

626 **Reduced Media Independent Interface**

627 **RMII**

628 a reduced signal count MAC to PHY interface, based on the IEEE Media Independent Interface (MII),
629 which was specified by the RMII Consortium (3Com Corporation; AMD Inc.; Bay Networks, Inc.;
630 Broadcom Corp.; National Semiconductor Corp.; and Texas Instruments Inc.)

631 **3.2.50**

632 **simple endpoint**

633 an MCTP endpoint that is not associated with either the functions of an MCTP bus owner or an MCTP
634 bridge.

635 **3.2.51**

636 **Transmission Unit**

637 refers to the size of the portion of the MCTP packet payload, which is the portion of the message body
638 carried in an MCTP packet.

639 **3.2.52**

640 **transport binding**

641 see [physical transport binding](#)

642 **3.2.53**

643 **Universally Unique Identifier**

644 **UUID**

645 refers to an identifier originally standardized by the Open Software Foundation (OSF) as part of the
646 Distributed Computing Environment (DCE)

647 Note 1 to entry: UUIDs are created using a set of algorithms that enables them to be independently
648 generated by different parties without requiring that the parties coordinate to ensure that generated IDs
649 do not overlap. In this specification, [RFC4122](#) is used as the base specification describing the format and
650 generation of UUIDs.

651 Note 2 to entry: Also, sometimes referred to as a globally unique identifier (GUID).

652 **4 Symbols and abbreviated terms**

653 The following symbols and abbreviations are used in this document.

654 4.1

655 **ACPI**

656 Advanced Configuration and Power Interface

657	4.2
658	ARP
659	Address Resolution Protocol
660	4.3
661	BCD
662	binary-coded decimal
663	4.4
664	BMC
665	baseboard management controller
666	4.5
667	CIM
668	Common Information Model
669	4.6
670	EID
671	endpoint identifier
672	4.7
673	FIFO
674	first-in first-out
675	4.8
676	GUID
677	Globally Unique Identifier
678	4.9
679	I²C
680	Inter-Integrated Circuit
681	4.10
682	IANA
683	Internet Assigned Numbers Authority
684	4.11
685	IP
686	Internet Protocol
687	4.12
688	IPMI
689	Intelligent platform management interface
690	4.13
691	ISO/IEC
692	International Organization for Standardization/International Engineering Consortium
693	4.14
694	MCTP
695	Management Component Transport Protocol

696 4.15
697 **MSB**
698 most significant byte

699 4.16
700 **PCle**
701 Peripheral Component Interconnect (PCI) Express

702 4.17
703 **PMCI**
704 Platform Management Component Intercommunications

705 4.18
706 **RMII**
707 Reduced Media Independent Interface

708 4.19
709 **SMBus**
710 System Management Bus

711 4.20
712 **TCP/IP**
713 Transmission Control Protocol/Internet Protocol

714 4.21
715 **USB**
716 Universal Serial Bus

717 4.22
718 **UUID**
719 Universally Unique Identifier

720 4.23
721 **VDM**
722 Vendor Defined Message

723 **5 Conventions**

724 **5.1 Overview**

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

726 **5.2 Byte ordering**

727 Unless otherwise specified, byte ordering of multi-byte numeric fields or bit fields is "Big Endian" (that is,
728 the lower byte offset holds the most significant byte, and higher offsets hold lesser significant bytes).

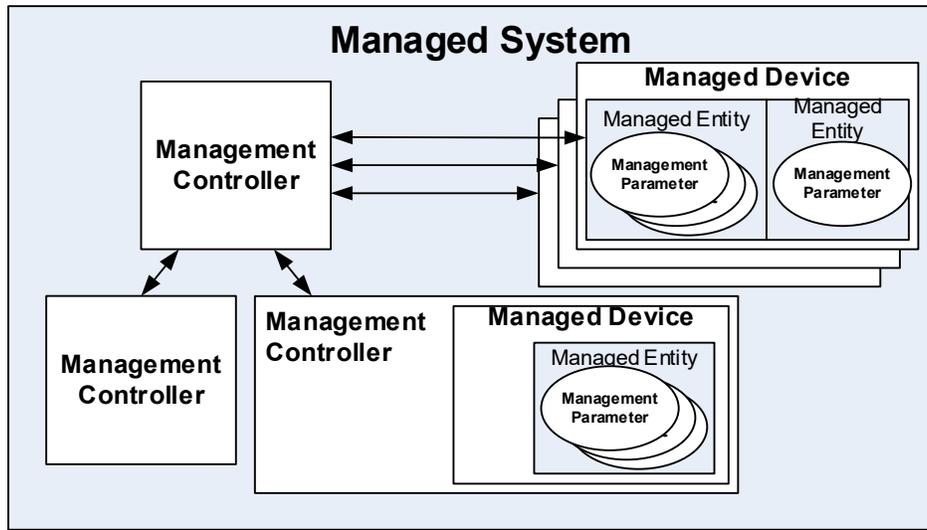
729 **5.3 Reserved fields**

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

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

734 **6 Management component relationships**

735 Figure 1 illustrates the relationship between devices, management controllers, managed devices, and
 736 managed entities, which are described in Clause 3.2.



737
 738 **Figure 1 – Management component relationships**

739 **7 MCTP overview**

740 This clause provides an overview of the main elements of MCTP. Additional overview information is
 741 available in the MCTP white paper, [DSP2016](#).

742 MCTP is a transport independent protocol that is used for intercommunication within an MCTP Network.
 743 An MCTP Network that consists of one or more physical transports that are used to transfer MCTP
 744 Packets between MCTP Endpoints. MCTP Transport Binding Specifications define how the MCTP
 745 protocol is implemented across a particular physical transport medium. For example, the DMTF has
 746 defined transport bindings for MCTP over [SMBus/12C](#) and MCTP over PCIe using PCIe Vendor Defined
 747 Messages (VDMs), and others.

748 An MCTP Endpoint is the terminus for MCTP communication. A physical device that supports MCTP may
 749 provide one or more MCTP Endpoints. Endpoints are addressed using a logical address called the
 750 Endpoint ID, or EID. EIDs in MCTP are analogous to IP Addresses in Internet Protocol networking. EIDs
 751 can be statically or dynamically allocated.

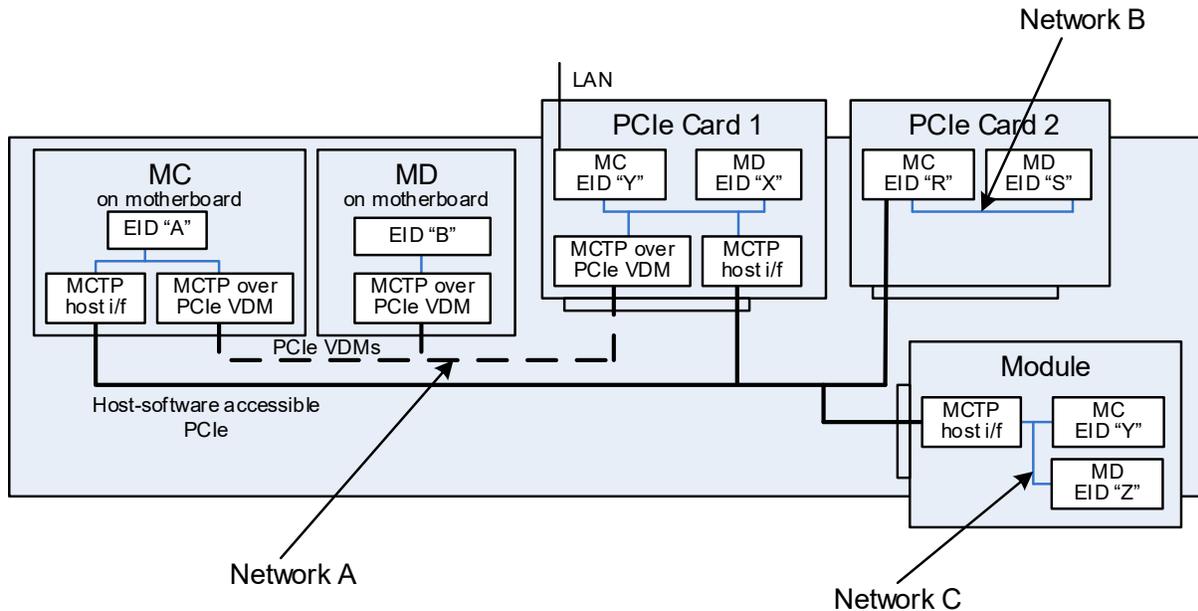
752 A system implementation can contain multiple MCTP Networks. Each MCTP Network has its own
 753 separate EID space. There is no coordination of EIDs between MCTP Networks. EIDs can overlap
 754 between MCTP Networks.

755 An MCTP Network may provide an MCTP Network ID that can be used to differentiate different MCTP
 756 Networks when more than one MCTP Network can be accessed by an entity such as system software.
 757 The Network ID is also used when an entity has more than one point of access to the MCTP Network. In

758 this case, the MCTP Network ID enables the entity to tell whether the access points provide access to the
 759 same MCTP Network or to different MCTP Networks.

760 The DMTF MCTP specifications also include the definition of transport bindings for MCTP host interfaces.
 761 MCTP host interfaces are used by software that runs locally on the host processors of the managed
 762 system to access an MCTP Network.

763



764

765

Figure 2 – MCTP networks

766 Figure 2 shows the different ways MCTP Networks can exist in a system. In this example, Network A
 767 connects a Management Controllers (MC) and managed devices (MD) on a motherboard, with devices on
 768 PCIe Card 1 using MCTP over PCIe Vendor Defined Messages. Note that there are two host interfaces
 769 (host i/f) on standard PCIe (host software accessible) that can be used by host software to access this
 770 particular network. This network thus requires an MCTP Network ID so that the host software can tell that
 771 the two host interfaces connect to the same MCTP Network.

772 Network B represents a network that is solely used for interconnecting devices within PCIe Card 2. This
 773 MCTP Network would typically not require an MCTP Network ID since it is not visible to host software or
 774 any other entity that would needs to differentiate Network B from another MCTP Network in the system.

775 Network C represents an MCTP Network on an add-in module. This network is separate from networks A
 776 and B but can accessed by host software through PCIe. Thus, this network requires a Network ID so that
 777 host software can differentiate that Network C is a different network than Network A.

778 MCTP Messages are comprised of one or more MCTP Packets. MCTP defines fields that support the
 779 assembly of received MCTP Packets into MCTP Messages and the disassembly of MCTP Messages into
 780 packets for transmission.

781 MCTP is designed to be able to transfer multiple Message Types in an interleaved manner using the
 782 same protocol. MCTP Message Types identified using a Message Type number. The use of the message
 783 type number is similar to a well-known port number in Internet Protocol. It identifies MCTP Messages that
 784 are all associated with a particular specification. This specification defines a Message Type for MCTP
 785 Control Messages that are used to initialize and maintain the MCTP Network. The DMTF has also defined

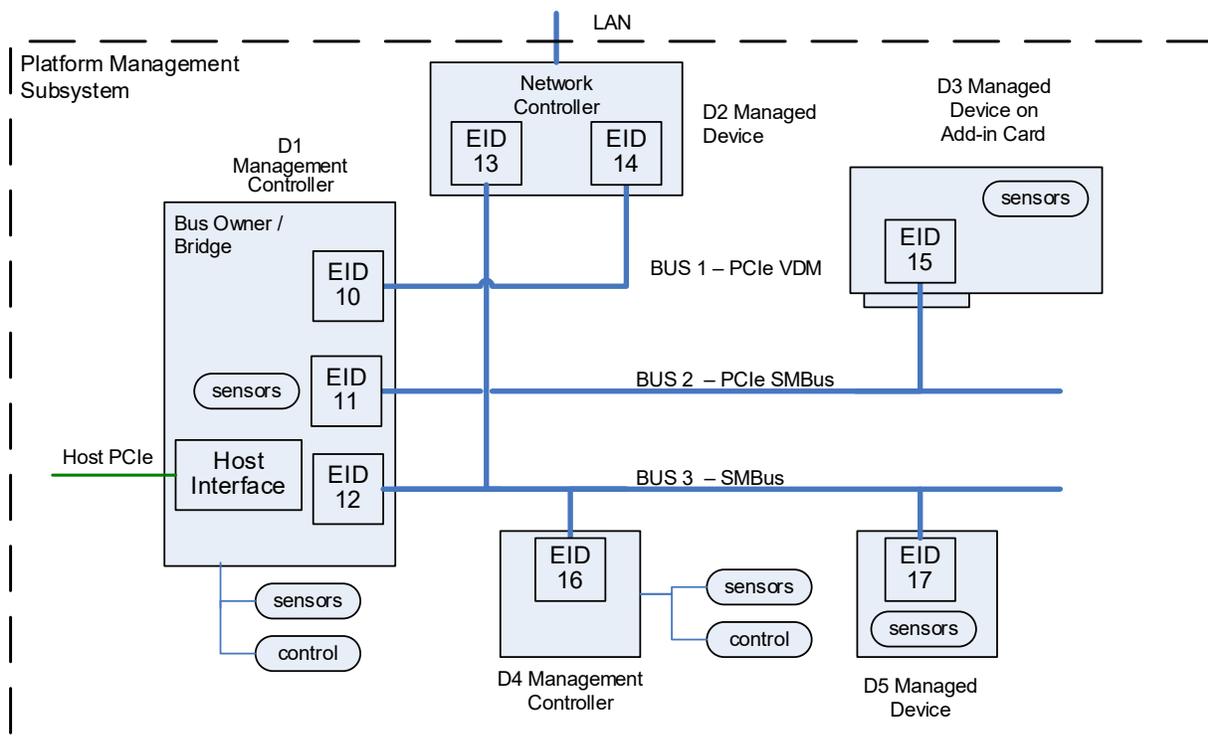
786 Message Types for use by the PMCI (Platform Management Communications Interconnect)
 787 specifications, Vendor-specific Messaging over MCTP, and so on. MCTP Message Type number
 788 assignments are provided in [DSP0239](#). [DSP0239](#) will be updated as new message types are defined in
 789 the future.

790 MCTP Control Messages use a request/response protocol. It is important to note that the base transport
 791 protocol defined by MCTP just defines a protocol for the transport of MCTP messages. Whether the
 792 message content is a request, a response, or something else is part of the particular Message Type
 793 definition.

794 In MCTP, a Bus is defined as a physical medium that shares a single physical address space. MCTP
 795 includes the definition of a function called the MCTP Bus Owner. The Bus Owner provides two main
 796 functions: It distributes EIDs to Endpoints when the MCTP implementation uses EIDs that are dynamically
 797 allocated, and it provides the way for an Endpoint to resolve an EID into the physical address used that is
 798 required to deliver a message to the target Endpoint.

799 Busses can be interconnected within an MCTP Network using MCTP Bridges to forward MCTP packets
 800 between busses. Bridges also handle the task of managing the difference in moving packets from one
 801 type of physical media to another, such as moving an MCTP packet between SMBus/I2C and PCIe
 802 Vendor Defined Messaging.

803 The following example illustrates how MCTP can be used within a hypothetical platform management
 804 subsystem implementation. More complex topologies, with multi-levels of bridges and greater numbers of
 805 busses and devices can be readily supported by MCTP as required.



806

807

Figure 3 – MCTP topology

808

809 **8 MCTP base protocol**

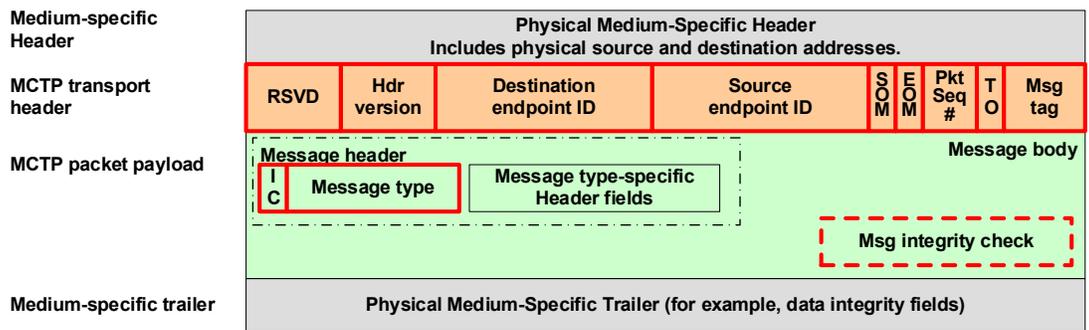
810 **8.1 Overview**

811 The MCTP base protocol defines the common fields for MCTP packets and messages and their usage.

812 Though there are medium-specific packet header fields and trailer fields, the fields for the base protocol
 813 are common for all media. These common fields support the routing and transport of messages between
 814 MCTP endpoints and the assembly and disassembly of large messages from and into multiple MCTP
 815 packets, respectively. The base protocol's common fields include a message type field that identifies what
 816 particular higher layer class of message is being carried using the MCTP base protocol.

817 **8.2 MCTP packet fields**

818 Figure 4 shows the fields that constitute a generic MCTP packet.



819

820 **Figure 4 – Generic message fields**

821 Table 1 defines the base protocol common fields.

822 **Table 1 – MCTP base protocol common fields**

Field Name	Field Size	Description
Medium-specific header	see description	This field represents the physical addressing and framing information that is used for transferring MCTP packets between devices on a particular physical medium. The size and type of any sub-fields or data within this field are defined by the corresponding transport binding specification for MCTP messaging on a given medium (for example, MCTP over SMBus/I2C, MCTP over PCIe Vendor Defined Messaging, and so on).
Medium-specific trailer	see description	This field represents any additional medium-specific trailer fields (if any) that are required for transferring MCTP packets between devices on a particular physical medium. A typical use of this field would be to hold per-packet data integrity fields (for example CRC, checksum, and so on) that would be specified for the particular medium.

Field Name	Field Size	Description
MCTP transport header	32 bits	The MCTP transport header is part of each MCTP packet and provides version and addressing information for the packet as well as flags and a "Message Tag" field that, in conjunction with the source EID, is used to identify packets that constitute an MCTP message. The MCTP transport header fields are common fields that are always present regardless of the physical medium over which MCTP is being used. Note: The positioning of the sub-fields of the MCTP transport header is allowed to vary based on the physical medium binding.
RSVD	4 bits	(Reserved) Reserved for future definition by the MCTP base specification.
Hdr version	4 bits	(Header version) Identifies the format, length, physical framing, and data integrity mechanism used to transfer the MCTP common fields in messages on a given physical medium. For this version of MCTP, this field shall be set to 0001b..
Destination endpoint ID	8 bits	The EID for the endpoint to receive the MCTP packet. A few EID values are reserved for specific routing. See Table 2 – Special endpoint IDs.
Source endpoint ID	8 bits	The EID of the originator of the MCTP packet. See Table 2 – Special endpoint IDs.
SOM	1 bit	(Start Of Message) Set to 1b if this packet is the first packet of a message.
EOM	1 bit	(End Of Message) Set to 1b if this packet is the last packet of a message.
Pkt Seq #	2 bits	(Packet sequence number) For messages that span multiple packets, the packet sequence number increments modulo 4 on each successive packet. This allows the receiver to detect up to three successive missing packets between the start and end of a message. Though the packet sequence number can be any value (0-3) if the SOM bit is set, it is recommended that it is an increment modulo 4 from the prior packet with an EOM bit set. After the SOM packet, the packet sequence number shall increment modulo 4 for each subsequent packet belonging to a given message up through the packet containing the EOM flag.
TO	1 bit	The TO (Tag Owner) bit identifies whether the message tag was originated by the endpoint that is the source of the message or by the endpoint that is the destination of the message. The Message Tag field is generated and tracked independently for each value of the Tag Owner bit. MCTP message types may overlay this bit with additional meaning, for example using it to differentiate between "request" messages and "response" messages. Set to 1b to indicate that the source of the message originated the message tag.
Msg tag	3 bits	(Message tag) Field that, along with the Source Endpoint IDs and the Tag Owner (TO) field, identifies a unique message at the MCTP transport level. Whether other elements, such as portions of the MCTP Message Data field, are also used for uniquely identifying instances or tracking retries of a message is dependent on the message type. A source endpoint is allowed to interleave packets from multiple messages to the same destination endpoint concurrently, provided that each of the messages has a unique message tag. When request/response message exchange is used and the Tag Owner (TO) bit is set to 1 in the request, a responder should return the same Message Tag with the Message Tag Owner bit cleared to 0 in the corresponding response Message. For messages that are split up into multiple packets, the Tag Owner (TO) and Message Tag bits remain the same for all packets from the SOM through the EOM.

Field Name	Field Size	Description
Message body	See description	The message body represents the payload of an MCTP message. The message body can span multiple MCTP packets.
IC	1 bit	(MCTP integrity check bit) Indicates whether the MCTP message is covered by an overall MCTP message payload integrity check. This field is required to be the most significant bit of the first byte of the message body in the first packet of a message along with the message type bits. 0b = No MCTP message integrity check 1b = MCTP message integrity check is present
Message type	7 bits	Defines the type of payload contained in the message data portion of the MCTP message. This field is required to be contained in the least-significant bits of the first byte of the message body in the first packet of a message. Like the fields in the MCTP transport header, the message type field is one of the common MCTP fields that are present independent of the transport over which MCTP is being used. Unlike the MCTP transport header, however, the message type field is only required to be present in the first packet of a particular MCTP message, whereas the MCTP transport header fields are present in every MCTP packet. See DSP0239 and Table 3 for information on message type values.
Message header	0 to M bytes	Additional header information associated with a particular message type, if any. This will typically only be contained in the first packet of a message, but a given message type definition can define header fields as required for any packet.
Message data	0 to N bytes	Data associated with the particular message type. Defined according to the specifications for the message type.
MCTP packet payload	See description	The packet payload is the portion of the message body that is carried in a given MCTP packet. The packet payload is limited according to the rules governing packet payload and transfer unit sizes. See 8.4, Packet payload and transmission unit sizes, for more information.
Msg integrity check	Message type-specific	(MCTP message integrity check) This field represents the optional presence of a message type-specific integrity check over the contents of the message body. If present, the Message integrity check field shall be carried in the last bytes of the message body. The particular message type definition will specify whether this is required, optional, or not to be used, the field size, and what algorithm is to be used to generate the field. The MCTP base protocol also does not specify whether this field is required on single packet messages (potentially dependent on transmission unit size) or is only required on multiple packet messages. Use of the Msg integrity check field is specific to the particular message type specification.

823 8.3 Special endpoint IDs

824 The following table lists EID values that are reserved or assigned to specific functions for MCTP.

825

Table 2 – Special endpoint IDs

Value	Description
Destination endpoint ID 0	Null Destination EID. This value indicates that the destination EID value is to be ignored and that only physical addressing is used to route the message to the destination on the given bus. This enables communication with devices that have not been assigned an EID. Because the physical addresses between buses are not guaranteed to be unique, MCTP does not support bridging messages with a null destination EID between different buses.

Value	Description
Source endpoint ID 0	Null Source EID. This value indicates a message is coming from an endpoint that is using physical addressing only. This would typically be used for messages that are delivered from an endpoint that has not been assigned an EID. Because the physical addresses between buses are not guaranteed to be unique, MCTP does not support bridging messages with a null source EID between different buses.
Endpoint IDs 1 through 7	Reserved for future definition.
Endpoint ID 0xFF	Broadcast EID. Reserved for use as a broadcast EID on a given bus. MCTP network-wide broadcasts are not supported. Primarily for use by the MCTP control message type.
All other values	Available for assignment and allocation to endpoints.

8.4 Packet payload and transmission unit sizes

8.4.1 Packet payload size

For MCTP, the size of a transmission unit is defined as the size of the packet payload that is carried in an MCTP packet.

8.4.2 Baseline transmission unit

The following are key information points regarding baseline transmission unit:

- The baseline transmission unit (minimum transmission unit) size for MCTP is 64 bytes.
- A message terminus that supports MCTP control messages shall always accept valid packets that have a transmission unit equal to or less than the baseline transmission unit. The message terminus is also allowed to support larger transmission units.
- The transmission unit of all packets in a given message shall be the same size, except for the transmission unit in the last packet (packet with EOM bit = 1b). Except for the last packet, this size shall be at least the baseline transmission unit size.
- The size of the transmission unit in the last packet shall be less than or equal to the transmission unit size used for the other packets (if any).
- If a transmission unit size larger than the baseline transmission unit is negotiated, the transmission unit of all packets shall be less than or equal to the negotiated transmission unit size. (The negotiation mechanism for larger transmission units between endpoints is message type-specific and physical medium-specific and is out of scope of this specification.)
- A given endpoint may negotiate additional restrictions on packet sizes for communication with another endpoint, as long as the requirements of this clause are met.
- All message types shall include support for being delivered using packets that have a transmission unit that is no larger than the baseline transmission unit. This is required to support bridging those messages in implementations where there are MCTP bridges that only support the baseline transmission unit.

8.5 Maximum message body sizes

The Message Body can span multiple packets. Limitations on message body sizes are message type-specific and are documented in the specifications for each message type.

854 8.6 Message assembly

855 The following fields (and *only* these fields) are collectively used to identify the packets that belong to a
856 given message for the purpose of message assembly on a particular destination endpoint when the
857 source EID is not 0.

- 858 • Msg Tag (Message Tag)
- 859 • TO (Tag Owner)
- 860 • Source Endpoint ID

861 The following fields (and *only* these fields) are collectively used to identify the packets that belong to a
862 given message for the purpose of message assembly on a particular destination endpoint when the
863 source EID is 0.

- 864 • Msg Tag (Message Tag)
- 865 • TO (Tag Owner)
- 866 • Source physical address

867 As described in 3.2, together these values identify the message terminus on the destination endpoint. For
868 a given message terminus, only one message assembly is allowed to be in process at a time.

869 8.7 Dropped packets

870 Individual packets are dropped (silently discarded) by an endpoint or physical layer under the following
871 conditions. These packets are discarded before being checked for acceptance or rejection for message
872 assembly. Therefore, these packets will *not* cause a message assembly to be started or terminated.

- 873 • **Unexpected "middle" packet or "end" packet**

874 A "middle" packet (SOM flag = 0 and EOM flag = 0) or "end" packet (SOM flag = 0 and EOM
875 flag = 1) for a multiple-packet message is received for a given message terminus without first
876 having received a corresponding "start" packet (where the "start" packet has SOM flag = 1 and
877 EOM flag = 0) for the message.

- 878 • **Bad packet data integrity or other physical layer error**

879 A packet is dropped at the physical data-link layer because a data integrity check on the packet
880 at that layer was invalid. Other possible physical layer errors may include framing errors, byte
881 alignment errors, packet sizes that do not meet the physical layer requirements, and so on.

- 882 • **Bad, unexpected, or expired message tag**

883 A message with TO bit = 0 was received, indicating that the destination endpoint was the
884 originator of the tag value, but the destination endpoint did not originate that value, or is no
885 longer expecting it. (MCTP bridges do not check message tag or TO bit values for messages
886 that are not addressed to the bridge's EID, or to the bridge's physical address if null-source or
887 destination-EID physical addressing is used.)

- 888 • **Un-routable EID**

889 An MCTP bridge receives an EID that the bridge is not able to route (for example, because the
890 bridge did not have a routing table entry for the given endpoint).

- 891 • **Bad header version**

892 The MCTP header version (Hdr Version) value is not a value that the endpoint supports.

893 • **Unsupported transmission unit**

894 The transmission unit size is not supported by the endpoint that is receiving the packet.

895 Individual packets should be dropped (silently discarded) by an endpoint under the following conditions.
896 These packets are discarded before being checked for acceptance or rejection for message assembly.
897 Therefore, these packets will *not* cause a message assembly to be started or terminated.

898 • **Unknown destination EID**

899 A packet is received at the physical address of a device which has a valid EID, but the
900 destination EID does not match the EID for the device.

901 • **Unsupported message type**

902 An MCTP endpoint receives a message type which is not supported by that endpoint.

903 **8.8 Starting message assembly**

904 Multiple-packet message assembly begins when the endpoint corresponding to the destination EID in the
905 packet receives a valid "start" packet (packet with SOM = 1b and EOM = 0b).

906 A packet with both SOM = 1b and EOM = 1b is considered to be a single-packet message, and is not
907 assembled per se.

908 Both multiple- and single-packet messages are subject to being terminated or dropped based on
909 conditions listed in the following clause.

910 **8.9 Terminating message assembly/dropped messages**

911 Message assembly is terminated at the destination endpoint and messages are accepted or dropped
912 under the following conditions:

913 • **Receipt of the "end" packet for the given message**

914 Receiving an "end" packet (packet with EOM = 1b) for a message that is in the process of being
915 assembled on a given message terminus will cause the message assembly to be completed
916 (provided that the message has not been terminated for any of the reasons listed below). This is
917 normal termination. The message is considered to be accepted at the MCTP base protocol
918 level.

919 • **Receipt of a new "start" packet**

920 Receiving a new "start" packet (packet with SOM = 1b) for a message to the same message
921 terminus as a message assembly already in progress will cause the message assembly in
922 process to be terminated. All data for the message assembly that was in progress is dropped.
923 The newly received start packet is not dropped, but instead it begins a new message assembly.
924 This is considered an error condition.

925 • **Timeout waiting for a packet**

926 Too much time occurred between packets of a given multiple-packet message. All data for the
927 message assembly that was in progress are dropped. This is considered an error condition. The
928 timeout interval, if specified, is specific to the transport binding specification. (A binding
929 specification may choose to not define a value for this timeout.)

- 930 • **Out-of-sequence packet sequence number**
- 931 For packets comprising a given multiple-packet message, the packet sequence number for the
932 most recently received packet is not a mod 4 increment of the previously received packet's
933 sequence number. All data for the message assembly that was in progress is dropped. This is
934 considered an error condition.
- 935 • **Incorrect transmission unit**
- 936 An implementation may terminate message assembly if it receives a "middle" packet (SOM =
937 0b and EOM = 0b) where the MCTP packet payload size does not match the MCTP packet
938 payload size for the start packet (SOM = 1b and EOM bit = 0b). This is considered an error
939 condition.
- 940 • **Bad message integrity check**
- 941 For single- or multiple-packet messages that use a message integrity check, a mismatch with
942 the message integrity check value can cause the message assembly to be terminated and the
943 entire message to be dropped, unless it is overridden by the specification for a particular
944 message type.
- 945 NOTE: The message integrity check is considered to be at the message-type level error condition rather
946 than an error at the MCTP base protocol level.

947 **8.10 Dropped messages**

- 948 An endpoint may drop a message if the message type is not supported by the endpoint. This can happen
949 in any one of the following ways:
- 950 • The endpoint can elect to not start message assembly upon detecting the invalid message type
951 in the first packet.
 - 952 • The endpoint can elect to terminate message assembly in process.
 - 953 • The endpoint can elect to drop the message after it has been assembled.

954 **8.11 MCTP versioning and message type support**

955 **8.11.1 MCTP version support**

956 There are three types of versioning information that can be retrieved using MCTP control messages:

- 957 • MCTP base specification version information
- 958 • MCTP packet header version information
- 959 • Message type version information

960 The version of the MCTP base specification that is supported by a given endpoint is obtained through the
961 Get MCTP Version Support command. This command can also be used to discover whether a particular
962 message type is supported on an endpoint, and if so, what versions of that message type are supported.

963 The Header Version field in MCTP packets identifies the media-specific formatting used for MCTP
964 packets. It can also indicate a level of current and backward compatibility with versions of the base
965 specification, as specified by the header version definition in each medium-specific transport binding
966 specification.

967 8.11.2 Compatibility with future versions of MCTP

968 An Endpoint may choose to support only certain versions of MCTP. The command structure along with
969 the Get MCTP Version Support command allows endpoints to detect and restrict the versions of MCTP
970 used by other communication endpoints. To support this, all endpoints on a given medium are required to
971 implement MCTP Version 1.0.x control commands or later 1.x Version for initialization and version
972 support discovery.
973

974 **8.12 MCTP message types**

975 Table 3 defines the values for the Message Type field for different message types transported through
 976 MCTP. The MCTP control message type is specified within this document. Baseline requirements for the
 977 Vendor Defined – PCI and Vendor Defined – IANA message types are also specified within this
 978 document. All other message types are specified in the [DSP0239](#) companion document to this
 979 specification.

980 NOTE: A device that supports a given message type is permitted to not support that message type equally across
 981 all buses that connect to the device.

982 **Table 3 – MCTP Message Types Used in this Specification**

Message Type	Message Type Code	Description
MCTP control	0x00	Messages used to support initialization and configuration of MCTP communication within an MCTP network. The messages and functions for this message type are defined within this specification.
Vendor Defined – PCI	0x7E	Message type used to support VDMs where the vendor is identified using a PCI-based vendor ID. The specification of the initial message header bytes for this message type is provided within this specification. Otherwise, the message body content is specified by the vendor, company, or organization identified by the given vendor ID.
Vendor Defined – IANA	0x7F	Message type used to support VDMs where the vendor is identified using an IANA-based vendor ID. (This format uses an "enterprise number" that is assigned and maintained by the Internet Assigned Numbers Authority, www.iana.org , as the means of identifying a particular vendor, company, or organization.) The specification of the initial message header bytes for this message type is provided within this specification. Otherwise, the message body content is specified by the vendor, company, or organization identified by the given vendor ID.

983 **8.13 Security**

984 The basic premise of MCTP is that higher layer protocols will fulfill security requirements (for example,
 985 confidentiality and authentication) for communication of management data. This means that the data
 986 models carried by MCTP shall fulfill the security requirements of a given management transaction. The
 987 MCTP protocol itself will not define any additional security mechanisms.

988 **8.14 Limitations**

989 MCTP has been optimized for communications that occur within a single computer system platform. It has
 990 not been designed to handle problems that can typically occur in a more generic inter-system networking
 991 environment. In particular, compared to networking protocols such as IP and TCP/IP, MCTP has the
 992 following limitations:

- 993 • MCTP has limited logical addressing. MCTP been optimized for the small number of endpoints
 994 that are expected to be utilized within the platform. The 8-bit range of EIDs is limited compared
 995 to the ranges available for IP addresses.
- 996 • MCTP assumes an MCTP network implementation that does not include loops. There is no
 997 mechanism defined in MCTP to detect or reconcile implementations that have connections that
 998 form routing loops.

- 999 • MCTP assumes a network topology where all packets belonging to a given message will be
1000 delivered through the same route (that is, MCTP does not generally support some packets for a
1001 message arriving by one route, while other packets for the message arrive by a different route).
- 1002 • MCTP does not support out-of-order packets for message assembly.
- 1003 • The MCTP base protocol does not address flow control or congestion control. These behaviors,
1004 if required, are specified at the physical transport binding level or at the message type or higher
1005 level.
- 1006 • MCTP is not specified to handle duplicate packets at the base protocol message assembly
1007 level. If a duplicate packet is received and passed on to MCTP message assembly, it can cause
1008 the entire message assembly to be terminated.
- 1009 NOTE: Transport bindings are not precluded from including mechanisms for handling duplicate packets
1010 at the physical transport level.

1011 **8.15 MCTP discovery and addressing**

1012 **8.15.1 Overview**

1013 This clause describes how MCTP endpoints, and their capabilities are discovered by one another, and
1014 how MCTP endpoints are provisioned with the addresses necessary for MCTP communication.

1015 MCTP discovery occurs over the course of several discrete, ordered steps:

- 1016 • Bus enumeration
- 1017 • Bus address assignment
- 1018 • MCTP capability discovery
- 1019 • Endpoint ID assignment
- 1020 • Distribution and use of routing information

1021 This clause gives an overview of the methods used for accomplishing each of these steps in various
1022 operational scenarios. Clause 13 gives details on the messages used to implement these operations.

1023 **8.15.2 Bus enumeration**

1024 This step represents existing bus enumeration. (The actions taken in this step are specific to a given
1025 medium.) Because enumeration of devices on the physical bus is medium-specific, this information is
1026 provided in the transport binding specification for the medium.

1027 **8.15.3 Bus address assignment**

1028 MCTP endpoints require a bus address that is unique to a given bus segment. This step deals with
1029 assignment of these addresses. Some bus types (such as PCIe) have built-in mechanisms to effectively
1030 deal with this. Others (such as SMBus/I2C) require some additional consideration. Because bus address
1031 assignment is medium-specific, this information is provided in the transport binding specification for the
1032 medium.

1033 **8.15.4 MCTP capability discovery**

1034 Capability discovery deals with the discovery of the characteristics of individual MCTP endpoints.
1035 Capabilities that can be discovered include what message types are supported by an endpoint and what
1036 message type versions are supported. See 8.11 for a description of the methods used to accomplish
1037 capability discovery.

1038 **8.15.5 Endpoint ID assignment**

1039 Endpoint IDs are system-wide unique IDs for identifying a specific MCTP endpoint. They can be
1040 dynamically assigned at system startup or hot-plug insertion. See 8.18 for a description of the methods
1041 used to accomplish EID assignment.

1042 **8.15.6 Distribution and use of routing information**

1043 Bridging-capable MCTP endpoints need routing information to identify the next hop to forward a message
1044 to its final destination. See clause 9 for a description of how routing information is conveyed between
1045 MCTP endpoints.

1046 **8.16 Devices with multiple media interfaces**

1047 MCTP fully supports management controllers or managed devices that have interfaces on more than one
1048 type of bus. For example, a device could have both a PCI Express (PCIe) and an SMBus/I2C interface. In
1049 this scenario, the device will typically have a different EID for each interface. (Bridges can include
1050 instantiations that have an endpoint shared across multiple interfaces; see 9.1.3 for more information.)

1051 This concept can be useful in different operational scenarios of the managed system. For example,
1052 typically a PCIe interface will be used during [ACPI "S0"](#) power states (when the system is fully powered
1053 up), which will provide significantly higher bandwidths, whereas the SMBus/I2C interface could be used
1054 for "S3–S5" low-power sleep states.

1055 The baseline transmission unit is specified to be common across all media, enabling packets to be routed
1056 between different media without requiring bridges to do intermediate assembly and disassembly
1057 operations to handle differences in packet payload sizes between different media.

1058 Devices that support multiple media interfaces shall meet the command requirements of this specification
1059 and the associated transport binding specification for each enabled interface. For a given message type,
1060 the device may implement the same message type –specific commands on all MCTP interfaces,
1061 regardless of the medium, unless otherwise specified by the message type specification.

1062 **8.17 Peer transactions**

1063 Endpoints can intercommunicate in a peer-to-peer manner using the physical addressing on a given bus.

1064 A special value for the EID is used in cases when the physical address is known, but the EID is not
1065 known. This capability is used primarily to support device discovery and EID assignment. A device that
1066 does not yet have an EID assignment is not addressed using an EID. Rather, the device gets its EID
1067 assigned using an MCTP control command, Set Endpoint ID, which uses physical addressing only.

1068 Similarly, depending on the transport binding, a device can also announce its presence by sending an
1069 MCTP message to a well-known physical address for the bus owner (for example, for PCIe VDM, this
1070 would be the root complex; for SMBus/I2C, the host slave address, and so on).

1071 It is important to note that in cases where two endpoints are on the same bus, they do not need to go
1072 through a bridge to communicate with each other. Devices use the Resolve Endpoint ID command to ask
1073 the bus owner what physical address should be used to route messages to a given EID. Depending on
1074 the bus implementation, the bus owner can either return the physical address of the bridge that the
1075 message should be delivered to, or it can return the physical address of the peer on the bus.

1076 **8.18 Endpoint ID assignment and endpoint ID pools**

1077 **8.18.1 Overview**

1078 MCTP EIDs are the system-wide unique IDs used by the MCTP infrastructure to address endpoints and
1079 for routing messages across multiple buses in the system. There is one EID assigned to a given physical
1080 or virtual address. Most managed devices or management controllers will connect to just a single bus and
1081 have a single EID. A non-bridge device that is connected to multiple different buses will have one EID for
1082 each bus it is attached to.

1083 Bus owners are MCTP devices that are responsible for issuing EIDs to devices on a bus segment. These
1084 EIDs come from a pool of EIDs maintained by the bus owner.

1085 With the exception of the topmost bus owner (see 8.18.2), a given bus owner's pools of EIDs are
1086 dynamically allocated at run-time by the bus owner of the bus above it in the hierarchy. Hot-plug devices
1087 shall have their EID pools dynamically allocated.

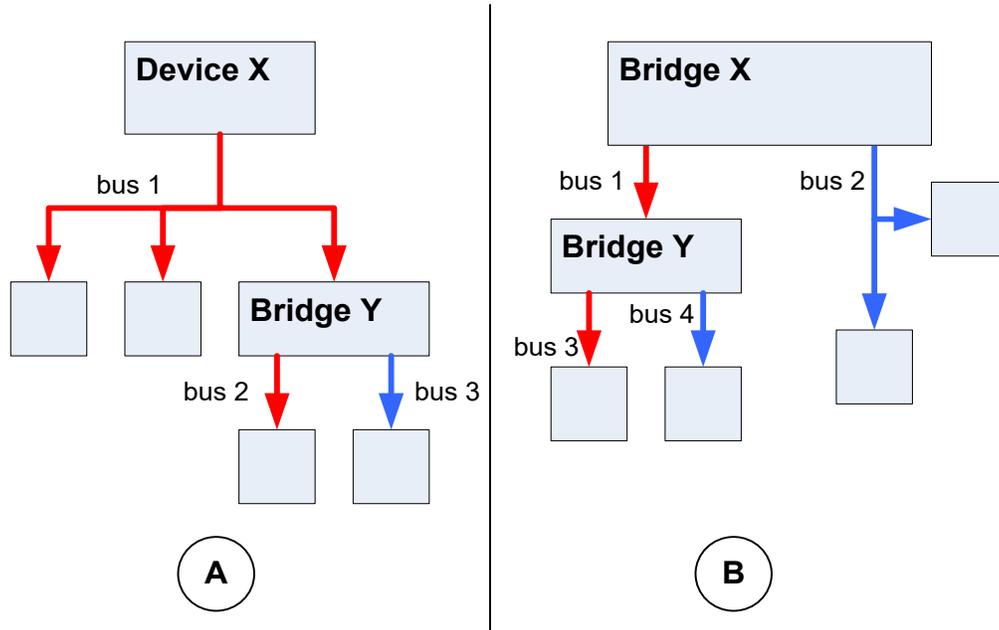
1088 Once EIDs are assigned to MCTP endpoints, it is necessary for MCTP devices involved in a transaction
1089 to understand something about the route a given message will traverse. Clause 9 describes how this
1090 routing information is shared among participants along a message's route.

1091 **8.18.2 Topmost bus owner**

1092 The topmost bus owner is the ultimate source of the EID pool from which all EIDs are drawn for a given
1093 MCTP network.

1094 1. This is illustrated in Figure 5, in which the arrows are used to identify the role of bus ownership. The
1095 arrows point outward from the bus owner for the particular bus and inward to a device that is "owned"
1096 on the bus.

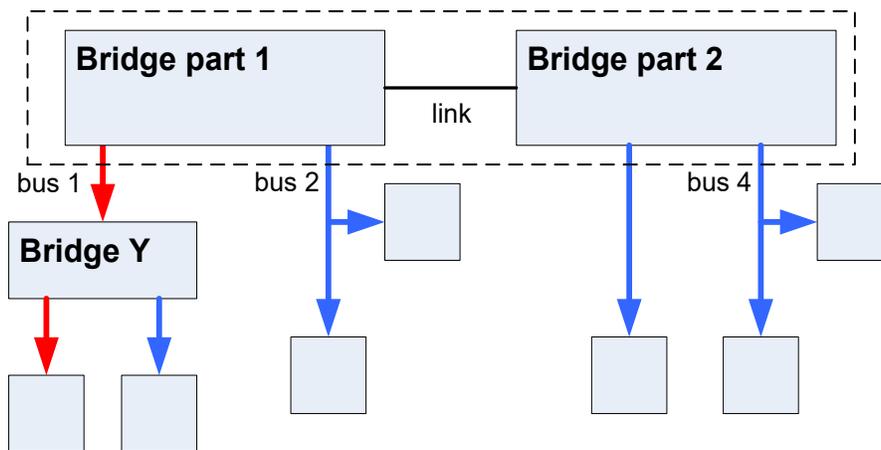
1097 In Figure 5, device X in diagram A and bridge X in diagram B are examples of topmost bus owners.
1098 Diagram A shows a device that connects to a single bus and is the topmost bus owner for the overall
1099 MCTP network. Diagram B shows that a bridge can simultaneously be the topmost bus owner, as well as
1100 the bus owner for more than one bus. The different colors represent examples of different media.



1101

1102

Figure 5 – Topmost bus owners



1103

1104

Figure 6 – Split bridge

1105 An implementation may need to split a bus owner or bridge across two physical devices. Such an
 1106 implementation shall include a mechanism (for example, a link as shown in Figure 6) that enables the two
 1107 parts to share a common routing table, or have individual copies of the routing table that are kept
 1108 synchronized. The definition of this mechanism is outside the scope of this specification.

1109 **8.18.3 Use of static EIDs and static EID pools**

1110 In general, the only device that will require a static (pre-configured default assigned non-zero value) EID
 1111 assignment will be the topmost bus owner. It needs a static EID because there is no other party to assign
 1112 it an EID through MCTP. Otherwise, all other devices will have their EIDs assigned to them by a bus
 1113 owner.

1114 The same principle applies if the device functions as an MCTP bridge. If the device is the highest device
1115 in the MCTP bus hierarchy, it will require a static pool of EIDs to be assigned to it as part of the system
1116 design. Otherwise, the device will be dynamically allocated pools of EIDs from a higher bus owner.

1117 An MCTP network implementation is allowed to use static EIDs for devices other than the topmost bus
1118 owner. Typically, this would only be done for very simple MCTP networks. Other key EID assignment
1119 considerations follow:

- 1120 • Endpoints that support the option of being configured for one or more static EIDs shall also
1121 support being configured to be dynamically assigned EIDs.
- 1122 • No mechanism is defined in the MCTP base specification for a bridge or bus owner to discover
1123 and incorporate a static EID into its routing information. Thus, a simple endpoint that is
1124 configured with a static EID shall also be used with a bus owner that is configured to support the
1125 static EIDs for the endpoint.
- 1126 • All bus owners/bridges in the hierarchy, from the topmost bus owner to the endpoint, shall have
1127 their routing configurable to support static EID routing information.
- 1128 • Although an endpoint that uses a static EID shall be used with a bus owner that supports static
1129 EIDs, the reverse is not true. A bus owner that uses static EIDs does not need to require that
1130 the devices on the buses it owns be configured with static EIDs.
- 1131 • How the configuration of static EIDs default value occurs is outside the scope of this
1132 specification.
- 1133 • No specified mechanism exists to "force" an override of a bridge's or bus owner's routing table
1134 entries for static EIDs. That is, commands such as Allocate Endpoint IDs and Routing
1135 Information Update only affect entries that are associated with dynamic EIDs.
- 1136 • MCTP does not define a mechanism for keeping routing tables updated if static EIDs are used
1137 with dynamic physical addresses. That is, static EIDs are not supported for use with dynamic
1138 physical addresses.
- 1139 • Bridges can have a mix of both static and dynamic EID pools. That is, the routing table can
1140 have both static and dynamic entries and can allocate from static and dynamic EID pools. Only
1141 the dynamic EID pools are given to the bridge by the bus owner using the Allocate Endpoint IDs
1142 command. There is no specification for how a static EID pool gets configured or how a bridge
1143 decides whether to give an endpoint an EID from a static or dynamically obtained EID pool.
1144 There is also no MCTP-defined mechanism to read the static EID pool setting from the bridge.
- 1145 • MCTP bridges and bus owners (except the topmost bus owner) are not required to include
1146 support for static EIDs.
- 1147 • MCTP does not define a mechanism for allocating EID pools that take static EID assignments
1148 into account. That is, a bridge cannot request a particular set of EIDs to be allocated to it.
- 1149 • MCTP bridges/bus owners may be configurable to use only static EIDs.

1150 **8.18.4 Use of static physical addresses**

1151 In many simple topologies, it is desirable to use devices that have statically configured physical
1152 addresses. This can simplify the implementation of the device. For example, an SMBus/I2C device that is
1153 not used in a hot-plug application would not need to support the SMBus address assignment (SMBus
1154 ARP) protocol. Fixed addresses can also aid in identifying the location and use of an MCTP device in a
1155 system. For example, if a system has two otherwise identical MCTP devices, a system vendor will know
1156 that the device at address "X" is the one at the front of the motherboard, and the device at address "Y" is
1157 at the back, because that is how they assigned the addresses when the system was designed.

1158 Therefore, MCTP transport bindings, such as for SMBus/I2C, are allowed to support devices being at
1159 static physical addresses without requiring the binding to define a mechanism that enables the bus owner
1160 to discover MCTP devices that are using static addresses.

1161 In this case, the bridge or bus owner shall have a-priori knowledge of the addresses of those devices to
1162 be able to assign EIDs to those devices and to support routing services for those devices. To support this
1163 requirement, the following requirements and recommendations are given to device vendors:

1164 • Devices that act as bus owners or bridges and are intended to support MCTP devices that use
1165 static physical addresses should provide a non-volatile configuration option that enables the
1166 system integrator to configure which device addresses are being used for devices on each bus
1167 that is owned by the bridge/bus owner.

1168 • The mechanism by which this non-volatile configuration occurs is specific to the device vendor.
1169 In many cases, the physical address information will be kept in some type of non-volatile
1170 storage that is associated with the device and gets loaded when the device is manufactured or
1171 when the device is integrated into a system. In other cases, this information may be coded into
1172 a firmware build for the device.

1173 8.18.5 Endpoint ID assignment process for bus owners/bridges

1174 The bus owner/bridge shall get its own EID assignment, and pools of EIDs, as follows. These steps only
1175 apply to bus owner/bridge devices that are not the topmost bus owner.

- 1176 • Bus owners/bridges shall be pre-configured with non-volatile information that identifies which
1177 buses they own. (How this configuration is accomplished is device/vendor specific and is
1178 outside the scope of this specification.)
- 1179 • The bus owner/bridge announces its presence on any buses *that it does not own* to get an EID
1180 assignment for that bus. The mechanism by which this announcement occurs is dependent on
1181 the particular physical transport binding and is defined as part of the binding specification.
- 1182 • The bus owner/bridge waits until it gets its own EID assignment for one of those buses through
1183 the Set Endpoint ID command.
- 1184 • The bus owner/bridge indicates the size of the EID pool it requires by returning that information
1185 in the response to the Set Endpoint ID command.
- 1186 • For each bus where the bus owner/bridge is itself an "owned" device, the bus owner/bridge will
1187 be offered a pool of EIDs by being sent an Allocate Endpoint IDs command from the bus owner.
- 1188 • The bus owner/bridge accepts allocations only from the bus of the "first" bus owner that gives it
1189 the allocation, as described in the Allocate Endpoint IDs command description in 8.10. If it gets
1190 allocations from other buses, they are rejected.
- 1191 • A bus owner/bridge can be allocated with additional EID pools, and may release some of the
1192 allocated EIDs which are not assigned as described in 8.18.8

1193 The bus owner can now begin to build a routing table for each of the buses that it owns and accept
1194 routing information update information. Refer to 9 for more information.

1195 8.18.6 Endpoint ID retention

1196 Devices should retain their EID assignments for as long as they are in their normal operating state.
1197 Asynchronous conditions, such as device errors, unexpected power loss, power state changes, resets,
1198 firmware updates, may cause a device to require a reassignment of its EID. Devices should retain their
1199 EID assignments across conditions where they may temporarily stop responding to commands over
1200 MCTP, such as during internal resets, error conditions, or configuration updates.

1201 8.18.7 Reclaiming EIDs from hot-plug devices

1202 Bridges will typically have a limited pool of EIDs from which to assign and allocate to devices. (This also
1203 applies when a single bus owner supports hot-plug devices.) It is important for bridges to reclaim EIDs so
1204 that when a device is removed, the EID can later be re-assigned when a device is plugged in. Otherwise,
1205 the EID pool could become depleted as devices are successively removed and added.

1206 EIDs for endpoints that use static addresses are not reclaimed.

1207 No mechanism is specified in the MCTP base protocol for detecting device removal when it occurs.
1208 Therefore, the general approach to detecting whether a device has been removed is to re-enumerate the
1209 bus when a new device is added, and an EID or EID pool is being assigned to that device.

1210 The following approach can be used to detect removed hot-plug devices: The bus owner/bridge can
1211 detect a removed device or devices by validating the EIDs that are presently allocated to endpoints that
1212 are directly on the bus and identifying which EIDs are missing. It can do this by attempting to access each
1213 endpoint that the bridge has listed in its routing table as being a device that is directly on the particular
1214 bus. Attempting to access each endpoint can be accomplished by issuing the Get Endpoint ID command
1215 to the physical address of each device and comparing the returned result to the existing entry in the
1216 routing table. If there is no response to the command, or if there is a mismatch with the existing routing
1217 information, the entry should be cleared, and the corresponding EID or EID range should be returned to
1218 the "pool" for re-assignment. The bus owner/bridge can then go through the normal steps for EID
1219 assignment.

1220 This approach should work for all physical transport bindings, because it keeps the "removed EID"
1221 detection processing separated from the address assignment process for the bus.

1222 In some cases, a hot-plug endpoint may temporarily go into a state where it does not respond to MCTP
1223 control messages. Depending on the medium, it is possible that when the endpoint comes back online, it
1224 does not request a new EID assignment but instead continues using the EID it had originally assigned. If
1225 this occurs while the bus owner is validating EIDs to see if any endpoints are no longer accessible, it is
1226 possible that the bus owner will assume that the endpoint was removed and reassign its EID to a newly
1227 inserted endpoint, unless other steps are taken:

- 1228 • The bus owner shall wait at least T_{RECLAIM} seconds before reassigning a given EID (where
1229 T_{RECLAIM} is specified in the physical transport binding specification for the medium used to
1230 access the endpoint).
- 1231 • Reclaimed EIDs shall only be reassigned after all unused EIDs in the EID pool have been
1232 assigned to endpoints. Optionally, additional robustness can be achieved if the bus owner
1233 maintains a short FIFO list of reclaimed EIDs (and their associated physical addresses) and
1234 allocates the older EIDs first.
- 1235 • A bus owner shall confirm that an endpoint has been removed by attempting to access it after
1236 T_{RECLAIM} has expired. It can do this by issuing a Get Endpoint ID command to the endpoint to
1237 verify that the endpoint is still non-responsive. It is recommended that this be done at least three
1238 times, with a delay of at least $1/2 * T_{\text{RECLAIM}}$ between tries if possible. If the endpoint continues
1239 to be non-responsive, it can be assumed that it is safe to return its EID to the pool of EIDs
1240 available for assignment.

1241 8.18.8 Reclaiming assigned EIDs and adding additional EIDs

1242 A bus owner may need to allocate more EIDs to a bridge, after it already allocated EID pools to other
1243 MCTP bridges. In this case, it can try to reclaim unassigned EIDs from bridges which were allocated with
1244 more EIDs than needed. Once unassigned EIDs are reclaimed, the bus owner can then assign these
1245 EIDs to the bridges which require additional EIDs.

1246 An MCTP Bridge which requires additional EIDs to be allocated, sends the [Request EID pool](#) to the bus
1247 owner which allocated its current pool of EIDs. Note that when an MCTP bridge has multiple endpoints, it
1248 shall only send the [Request EID pool](#) through the endpoint which was used to allocate its current pool of
1249 EIDs.

1250 If an MCTP bridge receives the [Request EID pool](#) from a downstream bridge and it has enough un-
1251 assigned EIDs it shall allocate the requested number of EIDs to the requesting downstream bridge.

1252 If an MCTP bridge receives the [Request EID pool](#) from a downstream bridge and it does not have enough
1253 un-assigned EIDs it should send another [Request EID pool](#) to the bus owner which allocated its current
1254 pool of EIDs, or alternatively it can reclaim unassigned EIDs from other downstream MCTP bridges if
1255 such exist.

1256 If the top-most MCTP bus owner receives the [Request EID pool](#) command and it does not have and
1257 cannot acquire the requested EIDs, it should respond with the available EIDs which it can allocate to the
1258 requesting MCTP bridge.

1259 To reclaim an EID pool from an MCTP bridge, the bus owner uses the [Release EID pool](#) control
1260 command to one or more bridges with EIDs which are not assigned. Adding pool(s) of EIDs to MCTP
1261 bridge(s) can be done using the [Allocate Endpoint IDs](#).

1262 **8.18.9 Managed devices with multiple endpoints on a given bus**

1263 When a managed device has multiple endpoints which are connected to the same bus, the device shall
1264 expose an MCTP bridge as described above and provide these multiple endpoints through that bridge.
1265 The internal endpoints will be exposed as being connected over a virtual bus of the same type as the
1266 bridge and will have virtual addresses assigned to them by the bridge. In this case the MCTP bridge only
1267 exposes the device endpoints and does not expose other external devices through that bridge.

1268 **8.18.10 Additional requirements for hot-plug endpoints**

1269 Devices that are hot-plug shall support the Get Endpoint UUID command. The purpose of this
1270 requirement is to provide a common mechanism for identifying when devices have been changed.

1271 Endpoints that go into states where they temporarily do not respond to MCTP control messages shall re-
1272 announce themselves and request a new EID assignment if they are "offline" for more than $T_{RECLAIM}$
1273 seconds, where $T_{RECLAIM}$ is specified in the physical transport binding specification for the medium used
1274 to access the endpoint.

1275 **8.18.11 Additional requirements for devices with multiple endpoints**

1276 A separate EID is utilized for each MCTP bus that a non-bridge device connects to. In many cases, it is
1277 desirable to be able to identify that the same device is accessible through multiple EIDs.

1278 If an endpoint has multiple physical interfaces (ports), the interfaces can be correlated to the device by
1279 using the MCTP Get Endpoint UUID command (see 13.6) to retrieve the unique system-wide identifier.

1280 Devices connected to multiple buses shall support the Get Endpoint UUID command for each endpoint
1281 and return a common UUID value across all the endpoints. This is to enable identifying EIDs as belonging
1282 to the same physical device.

1283 **8.19 Handling reassigned EIDs**

1284 Though unlikely, it is still possible that during the course of operation of an MCTP network, a particular
1285 EID could get reassigned from one endpoint to another. For example, this could occur if a newly hot-swap
1286 inserted endpoint device gets assigned an EID that was previously assigned to a device that was
1287 subsequently removed.

1288 Under this condition, it is possible that the endpoint could receive a message that was intended for the
1289 previously installed device. This is not considered an issue for MCTP control messages because the
1290 control messages are typically just used by bus owners and bridges for initializing and maintaining the
1291 MCTP network. The bus owners and bridges are aware of the EIDs they have assigned to endpoints and
1292 are thus intrinsically aware of any EID reassignment.

1293 Other endpoints, however, are not explicitly notified of the reassignment of EIDs. Therefore,
1294 communication that occurs directly from one endpoint to another is subject to the possibility that the EID
1295 could become assigned to a different device in the middle of communication. This shall be protected
1296 against by protocols specific to the message type being used for the communication.

1297 In general, the approach to protecting against this will be that other message types will require some kind
1298 of "session" to be established between the intercommunicating endpoints. By default, devices would not
1299 start up with an active session. Thus, if a new device is added and it gets a reassigned EID, it will not
1300 have an active session with the other device and the other device will detect this when it tries to
1301 communicate.

1302 The act of having a new EID assigned to an existing device should have the same effect. That is, if a
1303 device gets a new EID assignment, it would "close" any active sessions for other message types.

1304 The mechanism by which other message types would establish and track communication sessions
1305 between devices is not specified in this document. It is up to the specification of the particular message
1306 type.

1307 **9 MCTP bridging**

1308 **9.1.1 Overview**

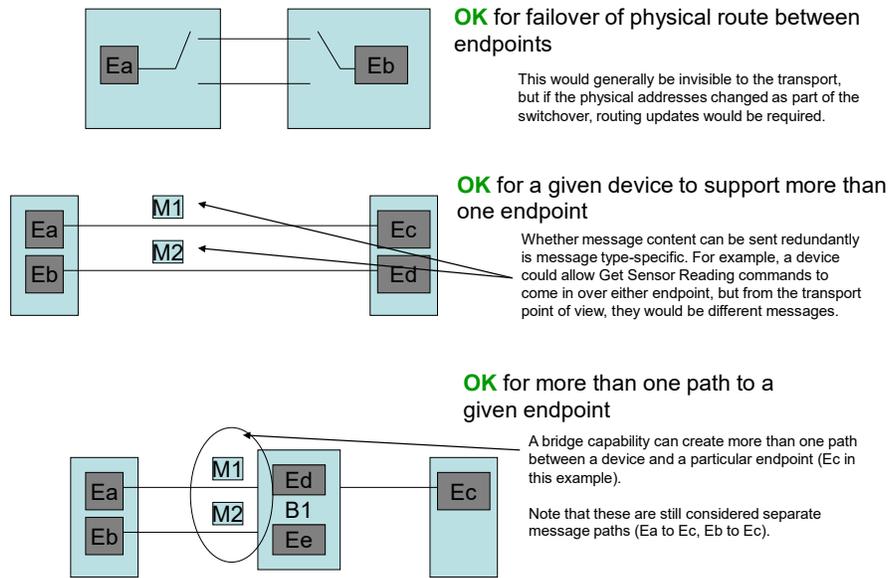
1309 One key capability provided by MCTP is its ability to route messages between multiple buses and
1310 between buses of different types. This clause describes how routing information is created, maintained,
1311 and used by MCTP bridges and MCTP endpoints. Keep the following key points in mind about MCTP
1312 bridges:

- 1313 • An MCTP bridge is responsible for routing MCTP packets between at least two buses.
- 1314 • An MCTP bridge is typically the bus owner for at least one of those buses.

1315 **9.1.2 Routing/bridging restrictions**

1316 Figure 7 and Figure 8 illustrate some of the supported and unsupported bridging topologies. As shown, it
1317 is acceptable for a given topology to have more than one path to get to a given EID. This can occur either
1318 because different media are used or because a redundant or failover communication path is desired in an
1319 implementation.

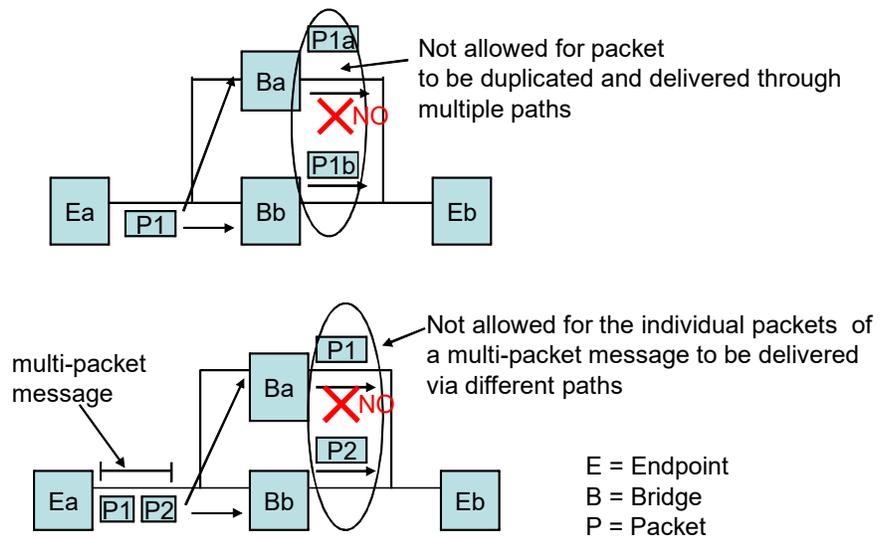
1320 A bridge shall not route or forward packets with a broadcast destination ID.



1321

1322

Figure 7 – Acceptable failover/redundant communication topologies



1323

1324

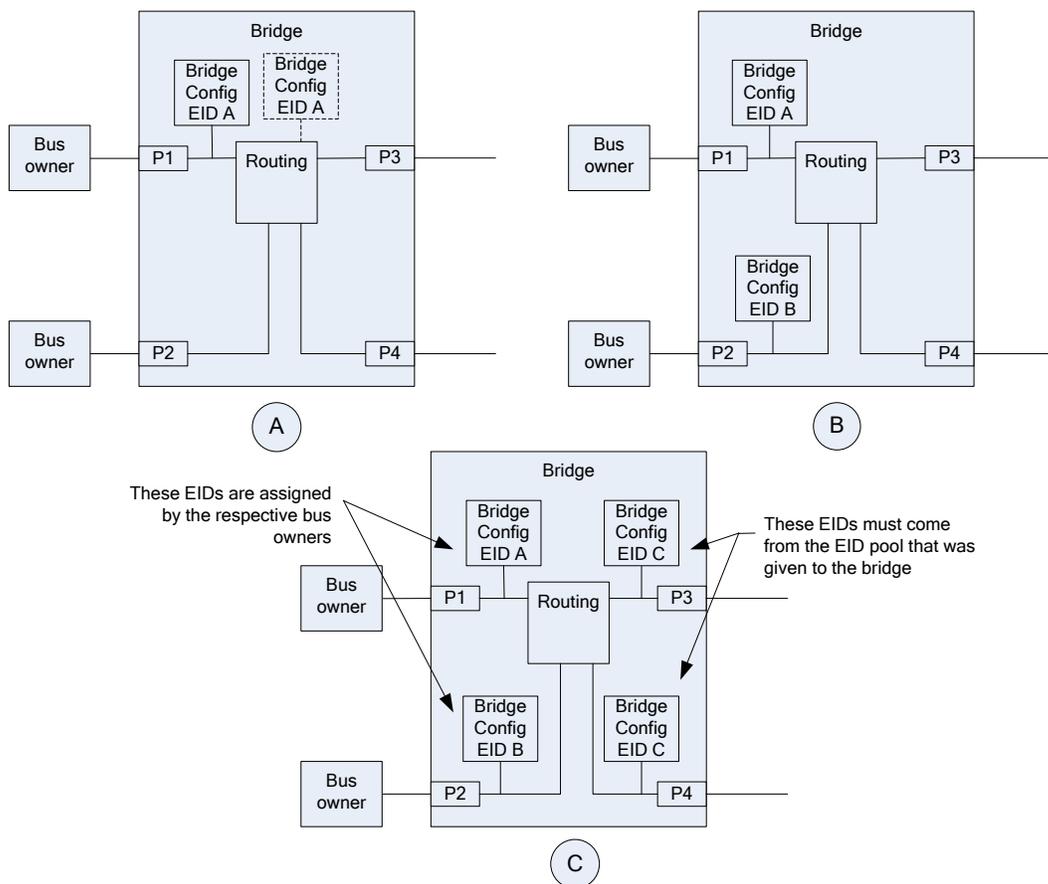
Figure 8 – Routing/bridging restrictions

1325 **9.1.3 EID options for MCTP bridges**

1326 An MCTP bridge that connects to multiple buses can have a single EID or multiple EIDs through which
 1327 the bridge's routing configuration and endpoint functionality can be accessed through MCTP control
 1328 commands. There are three general options:

- 1329 • The bridge uses a single MCTP endpoint
- 1330 • The bridge uses an MCTP endpoint for each bus that connects to a bus owner
- 1331 • The bridge uses an MCTP endpoint for every bus to which it connects

1332 Examples of these different options are shown in Figure 9, and more detailed information on the options
 1333 is provided following the figure.



1334
 1335 **Figure 9 – EID options for MCTP bridges**

1336 A bridge has one or more EID pools. To prevent issues with getting an EID pool allocation from multiple
 1337 bus owners, a bridge that is accessible through multiple EIDs will only accept EID pool allocation from the
 1338 first bus that allocation is received from using the Allocate Endpoint IDs command. This behavior is
 1339 described in more detail in the specification of the Allocate Endpoint IDs command.

1340 If necessary, the Get Endpoint UUID command can be used to correlate that EIDs belong to the same
 1341 MCTP bridge device. (This correlation is not required for normal initialization and operation of the MCTP
 1342 network, but it may be useful when debugging.)

1343 The following is a more detailed description of the different EID options for bridges:

1344 • **Single endpoint**

1345 A single endpoint is used to access the bridge's routing configuration and endpoint functionality.
1346 Referring to diagram (A) in Figure 9, an implementation may elect to either have the endpoint
1347 functionality be directly associated with a particular bus/port (for example, P1) or the
1348 functionality can be located on a "virtual bus" that is behind the routing function. In either case,
1349 the routing functionality ensures that the EID can be accessed through any of the buses to
1350 which the bridge connects.

1351 Although there is a single endpoint, the bridge shall report the need for EID assignment for that
1352 endpoint on each bus that is connected to a bus owner (for example, P1, P2). The multiple
1353 announcements provide a level of failover capability in the EID assignment process in case a
1354 particular bus owner becomes unavailable. The multiple announcements also help support a
1355 consistent EID assignment process across bus owners. To prevent issues with getting
1356 conflicting EID assignments from multiple bus owners, the bridge will only accept EID pool
1357 allocation from the first bus that an allocation is received from using the Set Endpoint ID
1358 command. This behavior is described in more detail in the specification of the Set Endpoint ID
1359 command. The bridge shall not report the need for EID assignment on any buses that the bridge
1360 itself owns.

1361 • **Endpoint for each bus connection to a bus owner**

1362 The bridge has one endpoint for each bus connected to a bus owner. This is shown as diagram
1363 (B) in Figure 9. There are no explicit endpoints associated with buses that are not connected to
1364 a bus owner (for example, the buses connected to ports P3 and P4, respectively.) Because of
1365 the way packet routing works, EID A and EID B can be accessed from any of the ports
1366 connected to the bridge. Thus, the bridge's configuration functionality may be accessed through
1367 multiple EIDs. Because a separate endpoint communication terminus is associated with each
1368 port (P1, P2), the bridge can accept an EID assignment for each bus independently.

1369 The bridge shall only report the need for EID assignment on buses that connect to a bus owner,
1370 and only for the particular MCTP control interface that is associated with the particular bus. For
1371 example, the bridge would announce the need for EID assignment for the interface associated
1372 with EID A only through P1, and the need for EID assignment for the interface associated with
1373 EID B only through P2. The bridge shall not report the need for EID assignment on any buses
1374 that the bridge itself owns.

1375 • **Endpoint for every bus connection**

1376 The bridge has one endpoint for each bus connected to it, as shown as diagram (C) in Figure 6.
1377 This includes buses that connect to bus owners (for example, P1, P2) and buses for which the
1378 bridge is the bus owner (for example, P3, P4). Because of the way packet routing works, any of
1379 these EIDs can be accessed from any of the ports connected to the bridge.

1380 Because a separate endpoint communication terminus is associated with each owned port (P1,
1381 P2), the bridge can accept an EID assignment for the bus owners of each bus independently.
1382 The EIDs associated with the buses that the bridge itself owns (for example, P3, P4) shall be
1383 taken out of the EID pool that is allocated to the bridge.

1384 The bridge shall only report the need for EID assignment on buses that connect to a bus owner,
1385 and only for the particular MCTP control interface that is associated with the particular bus. For
1386 example, the bridge would announce the need for EID assignment for the interface associated
1387 with EID A only through P1, and the need for EID assignment for the interface associated with
1388 EID B only through P2. The bridge shall not report the need for EID assignment on any buses
1389 that the bridge itself owns.

1390 9.1.4 Routing table

1391 An MCTP bridge maintains a routing table where each entry in the table associates either a single EID or
1392 a range of EIDs with a single physical address and bus ID for devices that are on buses that are directly
1393 connected to the bridge.

1394 If the device is a bridge, there will typically be a range of EIDs that are associated with the physical
1395 address of the bridge. There may also be an entry with a single EID for the bridge itself.

1396 9.1.5 Bridging process overview

1397 When a bridge receives an MCTP packet, the following process occurs:

- 1398 1) The bridge checks to see whether the destination EID in the packet matches or falls within the
1399 range of EIDs in the table.
- 1400 2) If the EID is for the bridge itself, the bridge internally consumes the packet.
- 1401 3) If there is a match with an entry in the routing table, the following steps happen:
 - 1402 • The bridge changes the physical addresses in the packet and reformats the medium-
1403 specific header and trailer fields as needed for the destination bus.
 - 1404 • The destination physical address from the source bus is replaced with the destination
1405 physical address for the destination bus obtained from the entry in the routing table.
 - 1406 • The bridge replaces the source physical address in the packet it received with the bridge's
1407 own physical address on the target bus. This is necessary to enable messages to be
1408 routed back to the originator.
 - 1409 • Packet-specific transport header and data integrity fields are updated as required by the
1410 particular transport binding.
- 1411 4) If there is no match, packets with EID values that are not in the routing table are silently
1412 discarded.

1413 9.1.6 Endpoint operation with bridging

1414 A bridge does not track the packet transmissions between endpoints. It simply takes packets that it
1415 receives and routes them on a per-packet basis based on the destination EID in the packet. It does not
1416 pay attention to message assembly or disassembly or message type-specific semantics, such as
1417 request/response semantics, for packets that it routes to other endpoints.

1418 Most simple MCTP endpoints will never need to know about bridges. Typically, another endpoint will
1419 initiate communication with them. The endpoint can then simply take the physical address and source
1420 EID information from the message and use that to send messages back to the message originator.

1421 An endpoint that needs to originate a "connection" to another MCTP endpoint does need to know what
1422 physical address should be used for messages to be delivered to that endpoint. To get this information, it
1423 needs to query the bus owner for it. An endpoint knows the physical address of the bus owner because it
1424 saved that information when it got its EID assignment.

1425 The Resolve Endpoint ID command requests a bus owner to return the physical address that is to be
1426 used to route packets to a given EID. (This is essentially the MCTP equivalent of the ARP protocol that is
1427 used to translate IP addresses to physical addresses.) The address that is returned in the Resolve
1428 Endpoint ID command response will either be the actual physical address for the device implementing the
1429 endpoint, or it will be the physical address for the bridge to be used to route packets to the desired
1430 endpoint.

1431 Because the physical address format is media-specific, the format of the physical address parameter is
 1432 documented in the specifications for the particular media-specific physical transport binding for MCTP (for
 1433 example, MCTP over SMBus/I2C, MCTP over PCIe Vendor Defined Messaging, and so on).

1434 If endpoint A has received a message from another endpoint B, it does not need to issue a Resolve
 1435 Endpoint ID command. Instead, it can extract the source EID and source physical address from the
 1436 earlier message from endpoint B, and then use that as the destination EID and destination physical
 1437 address for the message to Endpoint B.

1438 9.1.7 Routing table entries

1439 Each MCTP device that does bridging shall maintain a logical routing table. A bus owner shall also
 1440 typically maintain a routing table if more than one MCTP device is connected to the bus that it owns. The
 1441 routing table is required because the bus owner is also the party responsible for resolving EIDs to
 1442 physical addresses.

1443 The internal format that a device uses for organizing the routing table is implementation dependent. From
 1444 a logical point of view, each entry in a routing table will be comprised of at least three elements: An EID
 1445 range, a bus identifier, and a bus address. This is illustrated in Figure 10.

EID Range	Bus ID	Bus Address
-----------	--------	-------------

1446 **Figure 10 – Basic routing table entry fields**

1447 The *EID range* specifies the set of EIDs that can be reached through a particular bus address on a given
 1448 bus. Because the bus ID and bus address may correspond to a particular "port" on a bridge, it is possible
 1449 that there can be multiple non-contiguous ranges (multiple routing table entries) that have the same bus
 1450 ID/bus address pair route. EIDs and EID ranges can be categorized into three types: downstream,
 1451 upstream, and local. "Downstream" refers to EIDs that are associated with routing table entries that are
 1452 for buses that are owned by the bridge that is maintaining the routing table. "Upstream" refers to EIDs that
 1453 are associated with routing table entries that route to buses that are not owned by the bridge that is
 1454 maintaining the routing table.

1455 "Local" refers to the EIDs for routing table entries for endpoints that are on buses that are directly
 1456 connected to the bridge that is maintaining the routing table. A particular characteristic of entries for local
 1457 EIDs is that the Resolve Endpoint ID command is issued from the same bus that the endpoint is on. The
 1458 bridge/bus owner delivers the physical address for that endpoint rather than the physical address
 1459 associated with a routing function. This facilitates allowing endpoints on the same the bus to
 1460 communicate without having to go through an MCTP routing function.

1461 A routing table entry may not be "local" even if two endpoints are located on the same bus. An implementation may
 1462 require that different endpoints go through the routing function to intercommunicate even if the endpoints are part of
 1463 the same bus.

1464 The *bus ID* is an internal identifier that allows the MCTP device to identify the bus that correlates to this
 1465 route. MCTP does not require particular values to be used for identifying a given physical bus connection
 1466 on a device. However, this value will typically be a 0-based numeric value.

1467 EXAMPLE: A device that had three buses would typically identify them as buses "0", "1", and "2".

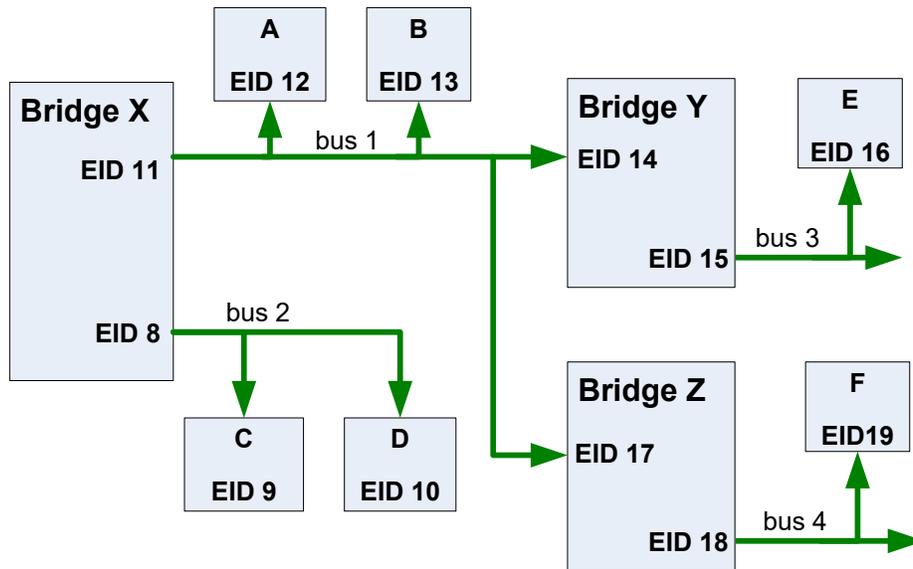
1468 The *bus address* is the physical address of a specific device on the bus through which the EIDs specified
 1469 in the *EID range* can be reached. This can either be the physical address corresponding to the
 1470 destination endpoint, or it can be the physical address of the next bridge in the path to the device. The
 1471 format of this address is specific to the particular physical medium and is defined by the physical medium
 1472 transport binding.

1473 **9.1.8 Routing table creation**

1474 **9.1.8.1 Overview**

1475 This clause illustrates the types of routing information that a bridge requires, and where the information
 1476 comes from. This clause also describes the steps that a bus owner shall use to convey that information
 1477 for a given bus.

1478 Figure 11 helps illustrate the steps that are required to completely establish the routing information
 1479 required by a bridge (bridge Y). The arrows in Figure 11 point outward from the bus owner and inward to
 1480 "owned" endpoints on the bus.



1481

1482 **Figure 11 – Routing table population**

1483 **9.1.8.2 Routing table population example**

1484 With reference to Figure 11, the following items describe the information that bridge Y will need for routing
 1485 messages in the example topology shown:

- 1486 • It needs a set of EIDs allocated to it to use for itself and to allocate to other devices (for
 1487 example, EIDs 14:16). These are allocated to it by the bus owner (bridge X).
- 1488 • It needs a routing table that has an entry that maps EID 16 to the physical address for device E
 1489 on bus 3.
- 1490 • It needs routing table entries for the local devices on bus 1, which are: bridge X (EID 11), device
 1491 A (EID 12), device B (EID 13), and bridge Z (EID 17), assuming that devices A and B are to be
 1492 reached by bridge Y without having to go through bridge X. This information shall be given to it
 1493 by the bus owner (bridge X).

- 1494 • It needs to know that EIDs 8:10 are accessed through bus owner/bridge X. Therefore, it needs a
1495 routing table entry that maps the EID range 9:10 to the physical address for bridge X on bus 1.
1496 This information shall also be given to it by the bus owner (bridge X).
- 1497 • It needs to know that EIDs 17:19 are accessed through bridge Z. Therefore, it needs a routing
1498 table entry that maps the EID range 17:19 to the physical address for bridge Z on bus 1.
1499 Because the bus owner (bridge X) allocated that range of EIDs to bridge Z in the first place, this
1500 information is also given to bridge Y by the bus owner (bridge X).

1501 9.1.8.3 Bus initialization example

1502 Starting with the description of what bridge Y requires, the following task list shows the steps that bridge
1503 X shall take to provide routing information for bus 1. Bridge X shall:

- 1504 1) Assign EIDs to devices A, B, C, D, bridge Y, and bridge Z. This is done using the Set Endpoint ID
1505 command. The response of the Set Endpoint ID command also indicates whether a device wants an
1506 additional pool of EIDs.
- 1507 2) Allocate EID pools to bridge Y and bridge Z. This is done using the Allocate Endpoint IDs command.
- 1508 3) Tell bridge Y the physical addresses and EIDs for devices A and B, bridge X (itself), and bridge Z on
1509 bus 1. This is done using the Routing Information Update command.
- 1510 4) Tell bridge Y that EIDs 18:19 are accessed through the physical address for bridge Z on bus 1. This
1511 is also done using the Routing Information Update command. (Steps 3 and 4 can be combined and
1512 covered with one instance of the command.)
- 1513 5) Tell bridge Z the physical addresses and EIDs for devices A and B, bridge X (itself), and bridge Y on
1514 bus 1. This is also done using the Routing Information Update command.
- 1515 6) Tell bridge Z that EIDs 15:16 are accessed through the physical address for bridge Y on bus 1. This
1516 is also done using the Routing Information Update command. (Steps 5 and 6 can be combined and
1517 covered with one instance of the command.)
- 1518 7) Tell bridge Y and bridge Z that EIDs 8:10 are accessed through bridge X on bus 1. This is also done
1519 using the Routing Information Update command. This step could also be combined with steps 3 and
1520 4 for bridge Y and steps 5 and 6 for bridge Z.

1521 9.1.9 Routing table updates responsibility for bus owners

1522 After it is initialized for all bridges, routing table information does not typically require updating during
1523 operation. However, updating may be required if a bridge is added as a hot-plug device. In this case,
1524 when the bridge is added to the system, it will trigger the need for the bus owner to assign it an EID,
1525 which will subsequently cause the request for EID pool allocations, and so on. At this time, the bus owner
1526 can simply elect to re-run the steps for bus initialization as described in 9.1.8.3.

1527 9.1.10 Consolidating routing table entries

1528 MCTP requires that when an EID pool is allocated to a device, each of the one or more ranges of EIDs is
1529 contiguous and does not include the EID for the bridge itself. Thus, a bridge can elect to consolidate
1530 routing table information into one entry when it recognizes that it has received an EID or EID range that is
1531 contiguous with an existing entry for the same physical address and bus. (The reason that EID allocation
1532 and routing information updates are not done as one range using the same command is because of the
1533 possibility that a device may have already received an allocation from a different bus owner.)

1534 **9.2 Bridge and routing table examples**

1535 **9.2.1 Overview**

1536 The following examples illustrate different bridge and MCTP network configurations and the
 1537 corresponding information that shall be retained by the bridge for MCTP packet routing and to support
 1538 commands such as Resolve Endpoint ID and Query Hop.

1539 The following clauses (including Table 4 through Table 6) illustrate possible topologies and ways to
 1540 organize the information that the bridge retains. Implementations may elect to organize and store the
 1541 same information in different ways. The important aspect of the examples is to show what information is
 1542 kept for each EID, to show what actions cause an entry to be created, and to show how an EID or EID
 1543 range can in some cases map to more than one physical address.

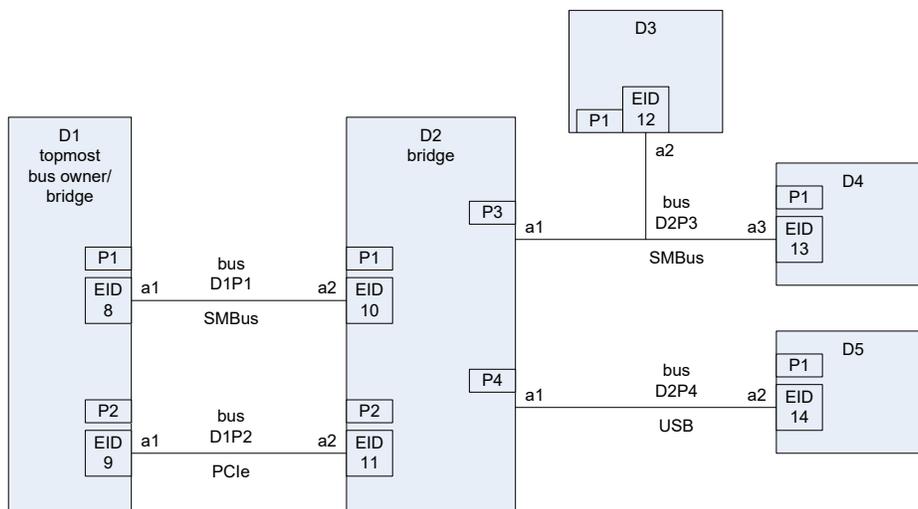
1544 The examples show a possible time order in which the entries of the table are created. Note that a given
 1545 implementation of the same example topology could have the entries populated in a different order. For
 1546 example, if there are two bus owners connected to a bridge, there is no fixed order that the bus owners
 1547 would be required to initialize a downstream bridge. Additionally, there is no requirement that bus owners
 1548 perform EID assignment or EID pool allocation in a particular order. One implementation may elect to
 1549 allocate EID pools to individual bridges right after it has assigned the bridge its EID. Another
 1550 implementation may elect to assign all the EIDs to devices first, and then allocate the EID pools to
 1551 bridges.

1552 **9.2.2 Example 1: Bridge D2 with an EID per "Owned" port**

1553 Figure 12 shows the routing table in a bridge (D2), where D2 has an EID associated with each bus
 1554 connected to a bus owner. In this example, D1 is not implementing any internal bridging between its P1
 1555 and P2. Consequently, EID9 cannot be reached by bridging through EID8 and vice versa (see Table 4).

1556 NOTE: If there was internal bridging, D1 would need to provide routing information that indicated that EID9 was
 1557 reachable by going through EID8 and vice versa. In this case, D1 would provide routing information that EID range
 1558 (EID8...EID9) would be accessed through D1P1a1 on SMBus and D1P2a1 on PCIe.

1559 **Key: D = device, P = port, a = physical address**



1560

1561

Figure 12 – Example 1 Routing topology

1562

Table 4 – Example 1 Routing table for D2

Time	EID	EID Access Port	Medium Type	Access Physical Address	Device/Entry Type	Entry Was Created and Populated By
	EID 10	P1	SMBus	D1P1a2	Bridge, Self	Self when EID was assigned by D1
	EID 11	P2	PCIe	D1P2a2	Bridge, Self	Self when EID was assigned by D1
	EID 12	P3	SMBus	D2P3a2	Endpoint	Self after D1 assigned EID pool (typically the entry will not be created until after the bridge D2 assigns EID 12 to D3)
	EID 13	P3	SMBus	D2P3a3	Endpoint	Self after D1 assigned EID pool (typically the entry will not be created until after the bridge D2 assigns EID 13 to D4)
	EID 14	P4	USB	D2P4a2	Endpoint	Self after D1 assigned EID pool (typically the entry will not be created until after the bridge D2 assigns EID 14 to D5)
	EID 8	P1	SMBus	D1P1a1	Bridge	D1 through Routing Information Update command
	EID 9	P2	PCIe	D1P2a1	Bridge	D1 through Routing Information Update command

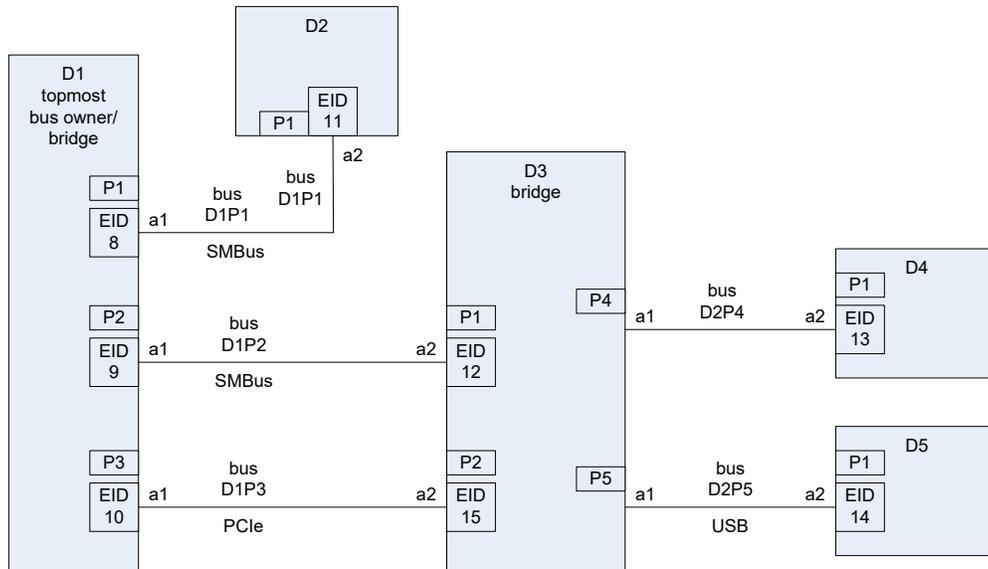
1563 **9.2.3 Example 2: Topmost bus owner D1**

1564 Figure 13 assumes the following conditions:

- 1565 • D1 assigns its internal EIDs first.
- 1566 • The buses are handled in the order D1P1, D1P2, D1P3.
- 1567 • D1 allocates the EID pool to bridges right after it has assigned the EID to the device.

1568 Similar to Example 1, this example assumes that there is no internal bridging within D1 between P1, P2,
 1569 and P3. This scenario is reflected in Table 5.

1570 Key: D = device, P = port, a = physical address



1571

Figure 13 – Example 2 Routing topology

1572

1573

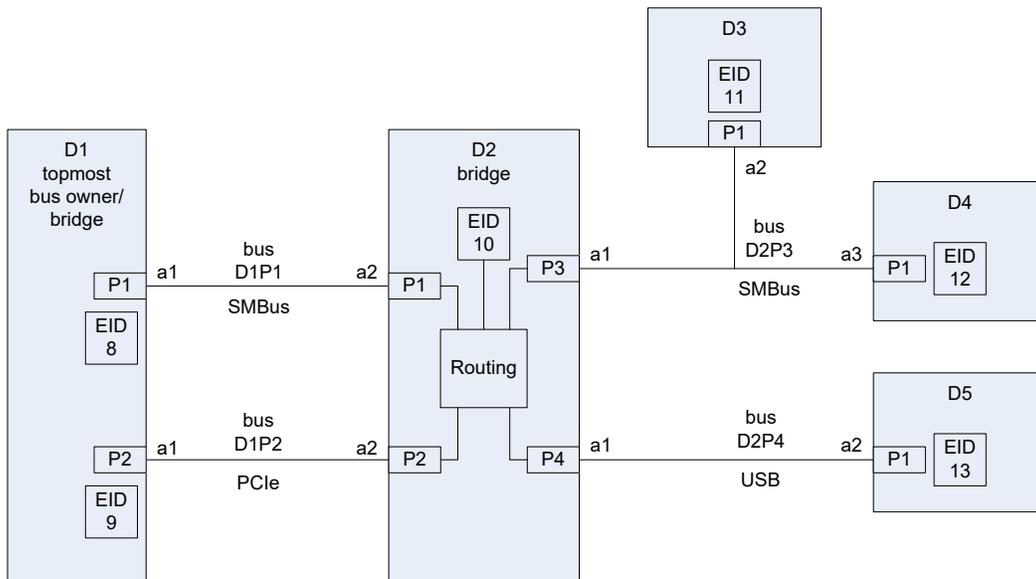
Table 5 – Example 2 Routing table for D1

EID	EID Access Port	Medium Type	Access Physical Address	Device/Entry Type	Entry Was Created and Populated By
EID 8	P1	SMBus	D1P1a1	Bridge, self	Self
EID 9	P2	SMBus	D1P2a1	Bridge, self	Self
EID 10	P3	PCIe	D1P3a1	Bridge, self	Self
EID 11	P1	SMBus	D1P1a2	Endpoint	Self upon assigning EID 11 to device D2
EID 12	P2	SMBus	D1P2a2	Bridge	Self upon assigning EID 12 to bridge D3
EID 13:14	P2	SMBus	D1P2a2	Bridge pool	Self upon assigning EID pool to bridge D3
EID 15	P3	PCIe	D1P3a2	Bridge	Self upon assigning EID 15 to bridge D3
EID 13:14	P3	PCIe	D1P3a2	Bridge pool	Self upon issuing an Allocate Endpoint IDs command and finding that bridge D3 already has an assigned pool, D1 creates this entry by extracting the EIDs for this entry from the response to the Allocate Endpoint IDs command

1574 **9.2.4 Example 3: Bridge D2 with single EID**

1575 Figure 14 assumes that bridge D2 has a single EID and gets its EID assignment and EID allocation
 1576 through bus D1P1 first, and that bus D1P2 later gets initialized. This scenario is reflected in Table 6.

1577 Key: D = device, P = port, a = physical address



1578

1579

Figure 14 – Example 3 Routing topology

1580

Table 6 – Example 3 Routing table for D2

Target EID	Target Endpoint Access Port	Target EID Access Physical Address	Device/Entry Type	Entry Was Created and Populated By
EID 10	P1	D1P1a2	Bridge, self	All four entries created by self (bridge) upon receiving initial EID assignment from D1 through P1
EID 10	P2	D1P2a2	Bridge, self	
EID 10	P3	D2P3a1	Bridge, self	
EID 10	P4	D2P4a1	Bridge, self	
EID 11	P3	D2P3a2	Endpoint	Self after D1 allocated EID pool (typically the entry will not be created until after the bridge D2 assigns EID 11 to D3)
EID 12	P3	D2P3a3	Endpoint	Self after D1 allocated EID pool (typically the entry will not be created until after the bridge D2 assigns EID 12 to D4)
EID 13	P3	D2P4a2	Endpoint	Self after D1 allocated EID pool (typically the entry will not be created until after the bridge D2 assigns EID 13 to D5)
EID 8:9	P1	D1P1a1	Bridge	D1 through Routing Information Update command
EID 8:9	P2	D1P2a1	Bridge	D1 through Routing Information Update command

1581 **9.2.5 Additional information tracked by bridges**

1582 In addition to the information required to route messages between different ports, a bridge has to track
 1583 information to handle MCTP control commands related to the configuration and operation of bridging
 1584 (shown in Table 7).

1585 **Table 7 – Additional information tracked by bridges**

What	Why
Which buses are connected to a bus owner	This information tells the bridge from which buses it should request EID assignment. This will typically be accomplished as a non-volatile configuration or hardware-strapping option for the bridge.
Which bus the bridge received its EID assignment through the Set Endpoint ID command	If the bridge uses a single EID that is shared across multiple "owned" buses, this information is used to track which bus the request came in on, so that the bridge can reject EID assignment requests from other buses.
Which bus it received the Routing Information Update command from for creating a particular routing table entry	This information is required so that if a future Routing Information Update command is received, the bridge will update only the entries corresponding to that bus.
Which bus it received its EID pool allocation from through the Allocate Endpoint IDs command	This information is used to track which bus the request came in on so that the bridge can reject EID pool allocations from other buses.
The physical medium and physical addressing format used for each port	This information is used to provide the correctly formatted response to commands such as Resolve Endpoint ID and for bridging MCTP packets between the different buses that the bridge supports. Because this is related to the physical ports and hardware of the bridge, this information will typically be "hard coded" into the bridge.

1586 **9.3 Endpoint ID resolution**

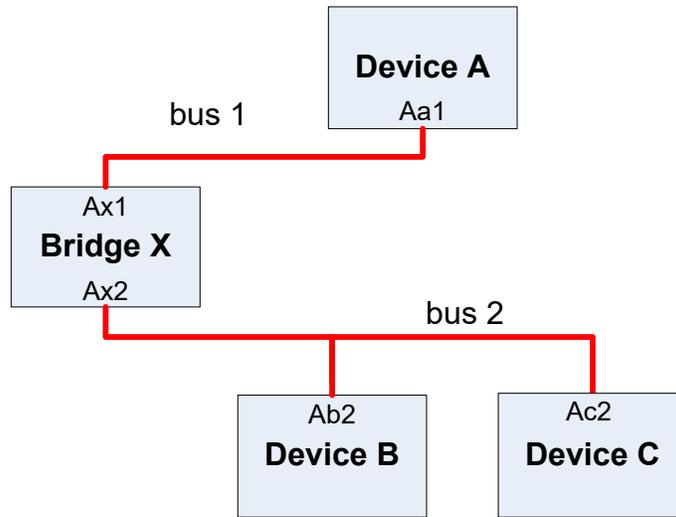
1587 **9.3.1 General**

1588 When a device uses the Resolve Endpoint ID command to request the resolution of a given endpoint to a
 1589 physical address, the bridge shall respond based on which bus the request came in on.

1590 For example, consider Figure 15. If device A wishes to get the physical address needed to send a
 1591 message to device C, it sends a Resolve Endpoint ID command to bus owner bridge X through address
 1592 Ax1. Because device A shall go through bridge X to get to device C, bridge X responds with its physical
 1593 address Ax1.

1594 When device B wishes to know the address to use to communicate with device C, it sends a Resolve
 1595 Endpoint ID request to bridge X through address Ax2. In this case, bridge X can respond by giving device
 1596 B the direct physical address of device C on bus 2, Ac2.

1597 Thus, the Resolve Endpoint ID command can return a different response based on the bus from which
 1598 the Resolve Endpoint ID command was received.



notation:

Ab2 = physical Address of device b on bus 2.

1599

1600

Figure 15 – Endpoint ID resolution

1601 **9.3.2 Resolving multiple paths**

1602 Cases can occur where there can be more than one possible path to a given EID. A likely scenario is
 1603 shown in Figure 16. In Figure 16, assume that the system topology supports cards that connect to either
 1604 SMBus, PCIe, or both. Bridge X is the bus owner for both buses.

1605 NOTE: This is a logical representation of MCTP buses. Physically, the buses may be formed of multiple physical
 1606 segments, as would be the case if one of the MCTP buses was built using PCIe.

1607 As shown, card C contains a bridge that connects to both buses. Thus, the device with EID 100 can be
 1608 reached either from bus 1 or bus 2.

1609 If device D wishes to send a message to EID 100, bridge X can choose to route that message either
 1610 through bus 1 or bus 2. MCTP does not have a requirement on how this is accomplished. The general
 1611 recommendation is that the bridge preferentially selects the faster available medium. In this example, that
 1612 would be PCIe.

1613 NOTE There are possible topologies where that simple rule is permitted to not yield the preferred path to a device.
 1614 However, in most common implementations in PC systems, this approach should be effective. A vendor making a
 1615 bridge device may consider providing configuration options to enable alternative policies.

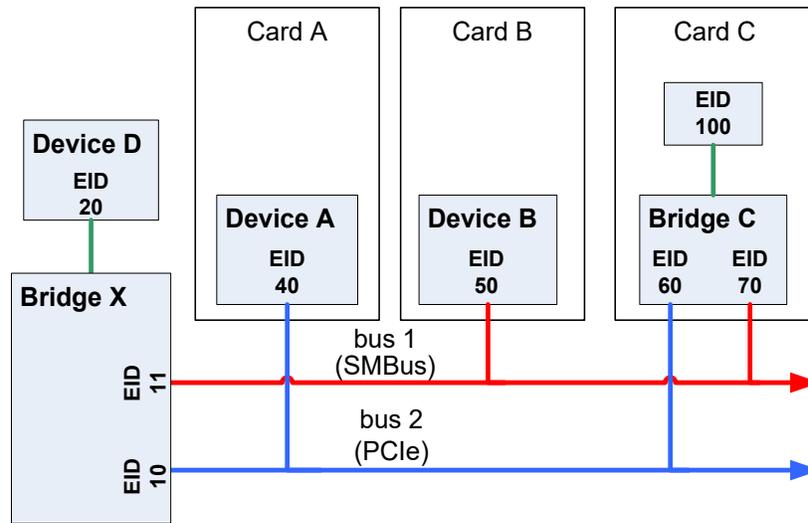


Figure 16 – Resolving multiple paths

1616

1617

1618 **9.4 Bridge and bus owner implementation recommendations**

1619 **9.4.1 Overview**

1620 This clause provides recommendations on EID pool and routing table sizes for devices that implement
1621 bridge and bus owner functionality.

1622 **9.4.2 Endpoint ID pool recommendations**

1623 The system design should seek to minimize the number of devices that need to allocate EID pools to hot-
1624 plug devices or add-in cards. If feasible, the system design should have all busses that support hot-plug
1625 devices/add-in cards owned by a single device.

1626 If only one device handles the hot-plug devices and add-in cards, it will be simpler for the system
1627 integrator to configure devices and allocate EID pools. Because any other bridges in the system that do
1628 not handle hot-plug devices only need to handle a fixed number of MCTP devices, it will be known at
1629 design time how large an EID pool will be required. The remaining number of EIDs can then simply be
1630 allocated to the single device that handles the hot-plug devices and add-in cards.

1631 To support this, it is recommended that devices that operate as bridges include a non-volatile
1632 configuration option that enables the system integrator to configure the size of the EID pool they request.

1633 **9.4.3 Routing table size recommendations**

1634 This clause provides some initial recommendations and approaches on how to determine what target
1635 routing table entry support to provide in a device.

1636 **• PCIe slots**

1637 To provide entries to support devices that plug into PCIe slots, assume that each slot may
1638 support both PCIe and SMBus endpoints and provide support for at least two endpoints per bus
1639 type.

1640 This means providing support for at least four directly connected endpoints per card. (Other
1641 endpoints may be behind bridges on the card, but this does not affect the routing table size for
1642 the bus owner.) This implies at least four routing table entries per PCIe slot. Thus, a device that

1643 was designed to support system implementations with eight PCIe slots should have support for
1644 32 routing table entries.

1645 • **Planar PCIe devices**

1646 In most PC systems, PCIe would be typically implemented as a single MCTP bus owned by a
1647 single device as the bus owner. Thus, the number of static devices should be proportional to the
1648 number of PCIe devices that are built into the motherboard.

1649 Typically, this is fewer than eight devices. Thus, it is recommended to support at least eight
1650 entries for static PCIe devices.

1651 • **Static SMBus/I2C MCTP devices**

1652 The routing table should also be sized to support an additional number of "static" devices on
1653 owned buses. At this time, it is considered unlikely that more than a few MCTP devices would
1654 be used on a given SMBus/I2C bus. Most devices would be non-intelligent sensor and I/O
1655 devices instead. Conservatively, it is recommended that at least four entries be provided for
1656 each SMBus/I2C bus that the device owns.

1657 Example 1: "client" capable device

1658	Four PCIe slots	→	16 routing table entries
1659	Two owned SMBus/I2C busses	→	+8 entries
1660	<u>Static PCIe device support</u>	→	<u>+8 entries</u>
1661			~32 entries or more

1662 Example 2: volume server capable

1663	Eight PCIe slots	→	32 routing table entries
1664	Four owned SMBus/I2C busses	→	+16 entries
1665	<u>Static PCIe device support</u>	→	<u>+8 entries</u>
1666			~56 entries or more

1667 **9.5 Path and transmission unit discovery**

1668 **9.5.1 Overview**

1669 The transmission unit is defined as the size of the MCTP packet payload that is supported for use in
1670 MCTP message assembly for a given message. The supported transmission unit sizes are allowed to
1671 vary on a per-message type basis.

1672 Intermediate bridges and physical media can limit the transmission unit sizes between endpoints.
1673 Therefore, the MCTP control protocol specifies a mechanism for discovering the transmission unit support
1674 for the path between endpoints when one or more bridges exist in the path between the endpoints.

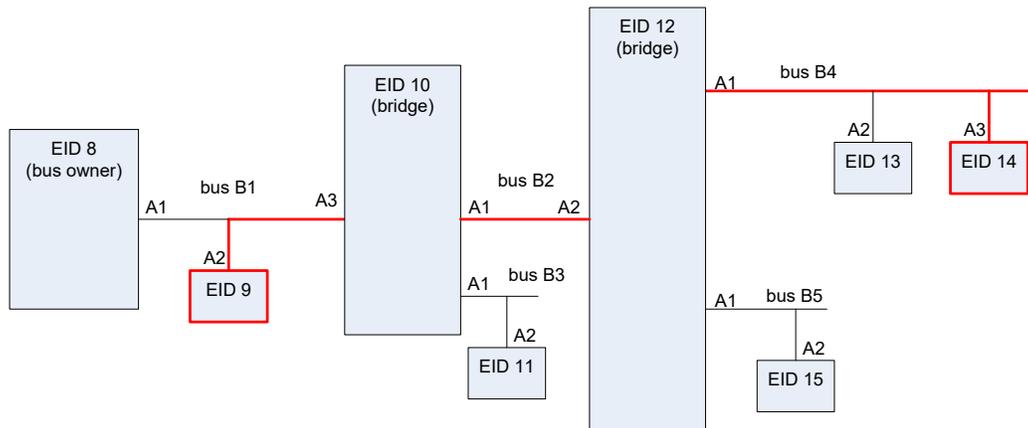
1675 The mechanism for path transmission unit discovery also enables the discovery of the bridges and
1676 number of "hops" that are used to route an MCTP packet from one endpoint to another.

1677 **9.5.2 Path transmission unit negotiation**

1678 The MCTP control protocol only specifies how to discover what the path transmission unit size is for the
1679 path between endpoints. The MCTP control protocol does not specify a generic mechanism for
1680 discovering what transmission unit sizes a particular endpoint supports for a given message type.
1681 Discovery and negotiation of transmission unit sizes for endpoints, if supported, is specified by the
1682 definition of the particular message type.

1683 **9.5.3 Path transmission unit discovery process overview**

1684 This clause describes the process used for path transmission unit discovery. The discovery process
 1685 described here is designed to enable one endpoint to discover the path and transmission unit support for
 1686 accessing a particular "target" endpoint. It does not define a general mechanism for enabling an endpoint
 1687 to discover the path between any two arbitrary endpoints. For example, referring to
 1688 Figure 17, the process defines a way for the endpoint at EID 9 to discover the path/transmission unit
 1689 support on the route to endpoint at EID 14, but this process does not define a process for EID 9 to
 1690 discover the path/transmission unit support between EID 11 and EID 14.



1691

1692

Figure 17 – Example path routing topology

1693 The following example provides an overview of the path/transmission unit discovery process. The
 1694 example presumes that the MCTP network has already been initialized. Referring to
 1695 Figure 17, the endpoint with EID 9 wishes to discover the path used to access the endpoint with EID 14.
 1696 This discovery is accomplished using just two commands, Resolve Endpoint ID and Query Hop, as
 1697 follows:

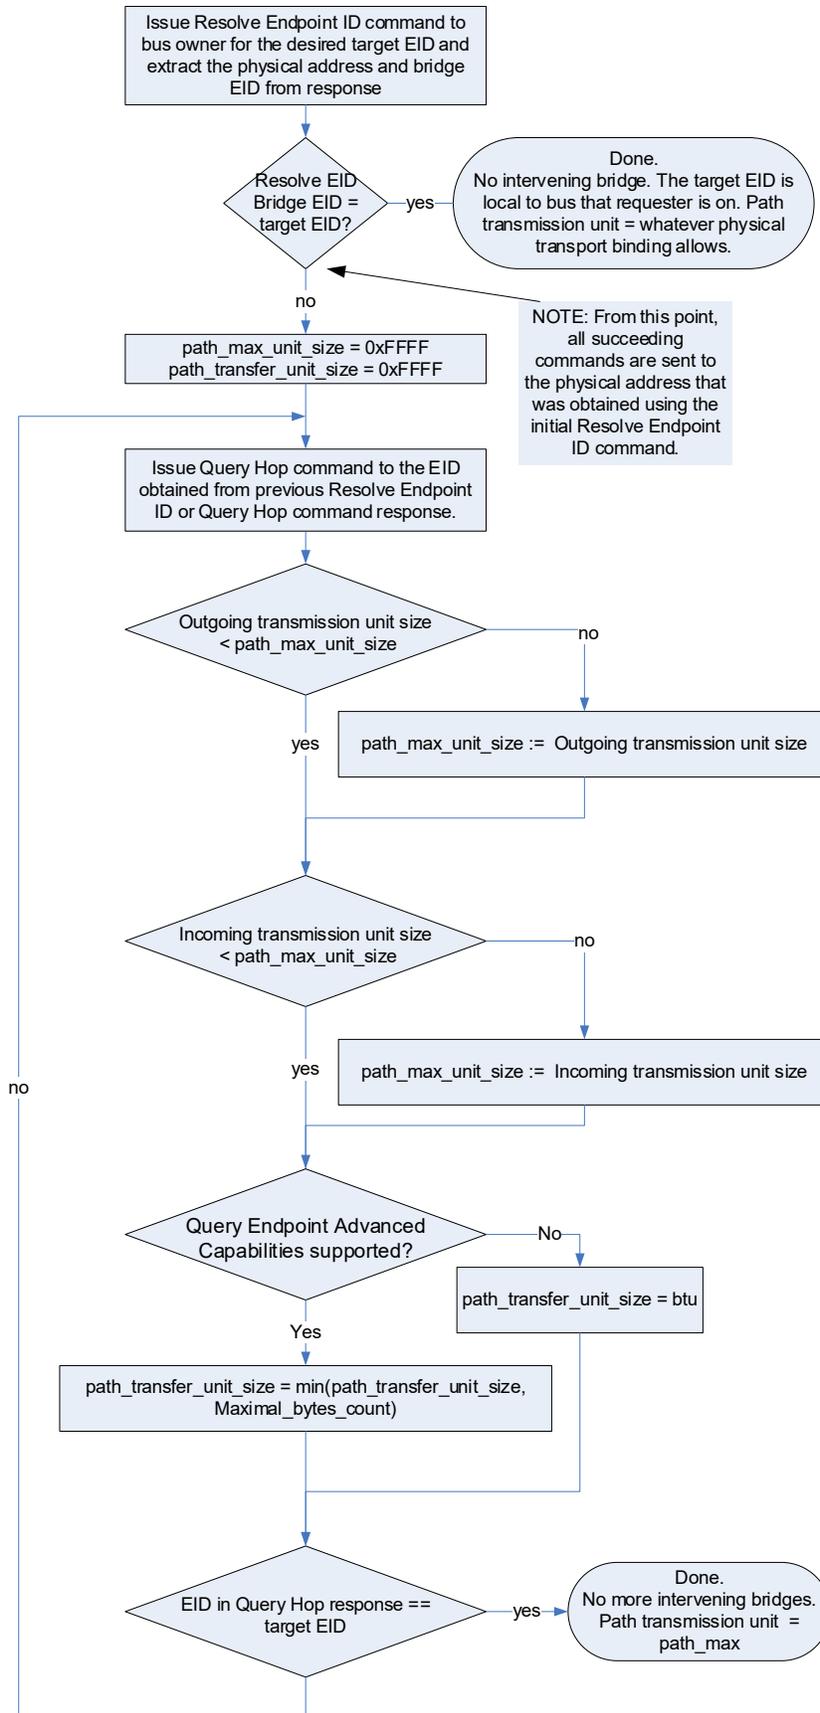
- 1698 1) EID 9 first issues a Resolve Endpoint ID command to the bus owner, EID 8, with EID 14 as the EID
 1699 to resolve.
- 1700 2) EID 8 returns the physical address and EID of the bridge, EID 10 in the Resolve Endpoint ID
 1701 command response.
- 1702 3) EID 9 queries the bridge, EID 10, using a Query Hop command with EID 14 (the "target" EID) as the
 1703 request parameter. Note that EID 9 does not need to do another Resolve Endpoint ID
 1704 command because it already received the physical address of EID 10 from the original Resolve Endpoint ID
 1705 command.
- 1706 4) Bridge EID 10 responds to the Query Hop command by returning EID 12, which is the EID of the
 1707 next bridge required to access EID 14. The bridge EID 10 also returns the transmission unit support
 1708 that it offers for routing to the target EID.
- 1709 5) EID 9 then sends a Query Hop command to the bridge at EID 12. Note that EID 9 does not need to
 1710 do another Resolve Endpoint ID command because it already received the physical address of EID
 1711 12 from the original Resolve Endpoint ID command.
- 1712 6) Bridge EID 12 responds to the Query Hop command by returning EID 14, which, because it is the
 1713 EID of the target endpoint, tells EID 9 that bridge EID 12 was the last "hop" in the path to EID 6. The
 1714 bridge EID 12 also returns the transmission unit support that it offers for routing to the target EID.

- 1715 7) At this point, the bridges in the path to EID 14 have subsequently been discovered and their
1716 respective transmission unit support returned. The effective transmission unit support for the path to
1717 EID 14 will be the lesser of the transmission unit support values returned by the two bridges.

1718 **9.5.4 Path transmission unit discovery process flowchart**

1719 The following flowchart (Figure 18) shows a generic algorithm for discovering the bridges in the path from
1720 one endpoint to a given target endpoint and the path transmission unit support. The flowchart has been
1721 intentionally simplified. Note that while the Query Hop command actually supports returning separate
1722 transmission unit sizes for the transmit and receive paths, the flowchart is simplified for illustration
1723 purposes and just refers to a single transmission unit for both transmit and receive.

1724 Additionally, Figure 18 does not show any explicit steps for error handling nor the process of handling
1725 command retries. In general, errors are most likely due to either an invalid EID being sent to the bridge
1726 (perhaps due to a programming error at the requester) or the EID not being present in the bridge's routing
1727 table. The latter condition could occur under normal operation if the requester did not realize that a
1728 routing table update had occurred because of a hot-plug update, for example. This error condition would
1729 be indicated by the bridge responding with an `ERROR_INVALID_DATA` completion code.



1731 **Figure 18 – Path transmission unit discovery flowchart**

1732 **9.6 Path transmission unit requirements for bridges**

1733 An MCTP bridge routes packets between different buses, but it does not typically interpret the packet
1734 payload contents nor does it do assembly of those packets. Exceptions to this are when the bridge is
1735 handling packets addressed to its own EID, receives a Broadcast EID, and if the bridge supports different
1736 transmission units based on message type. See Table 32 for more information.

1737 **10 Rate Limiting**

1738 **10.1.1 Methods**

1739 Some MCTP bindings provide a significant transfer rate that may not be sustainable by the MCTP
1740 message receiver. It is not always possible to use the native flow control mechanisms of the medium,
1741 since they may be shared with other traffic. In order to help address this problem, Endpoints may support
1742 the following specified MCTP Rate Limiting method.

1743 Note: The PCIe binding is a typical example of this issue. PCIe provides significantly more bandwidth
1744 than most MCTP endpoints can consume. PCIe credits cannot be used to throttle the MCTP traffic, since
1745 this would throttle all PCIe traffic (MCTP and non-MCTP) to the device. Thus, an alternative Rate Limiting
1746 mechanism is needed. Rate limiting is performed independently in each direction and is not required to be
1747 symmetric. Rate limiting can be set for one-direction only, for both directions or not be set at all.

1748 The MCTP rate limit mechanism allows an endpoint on a specific medium to:

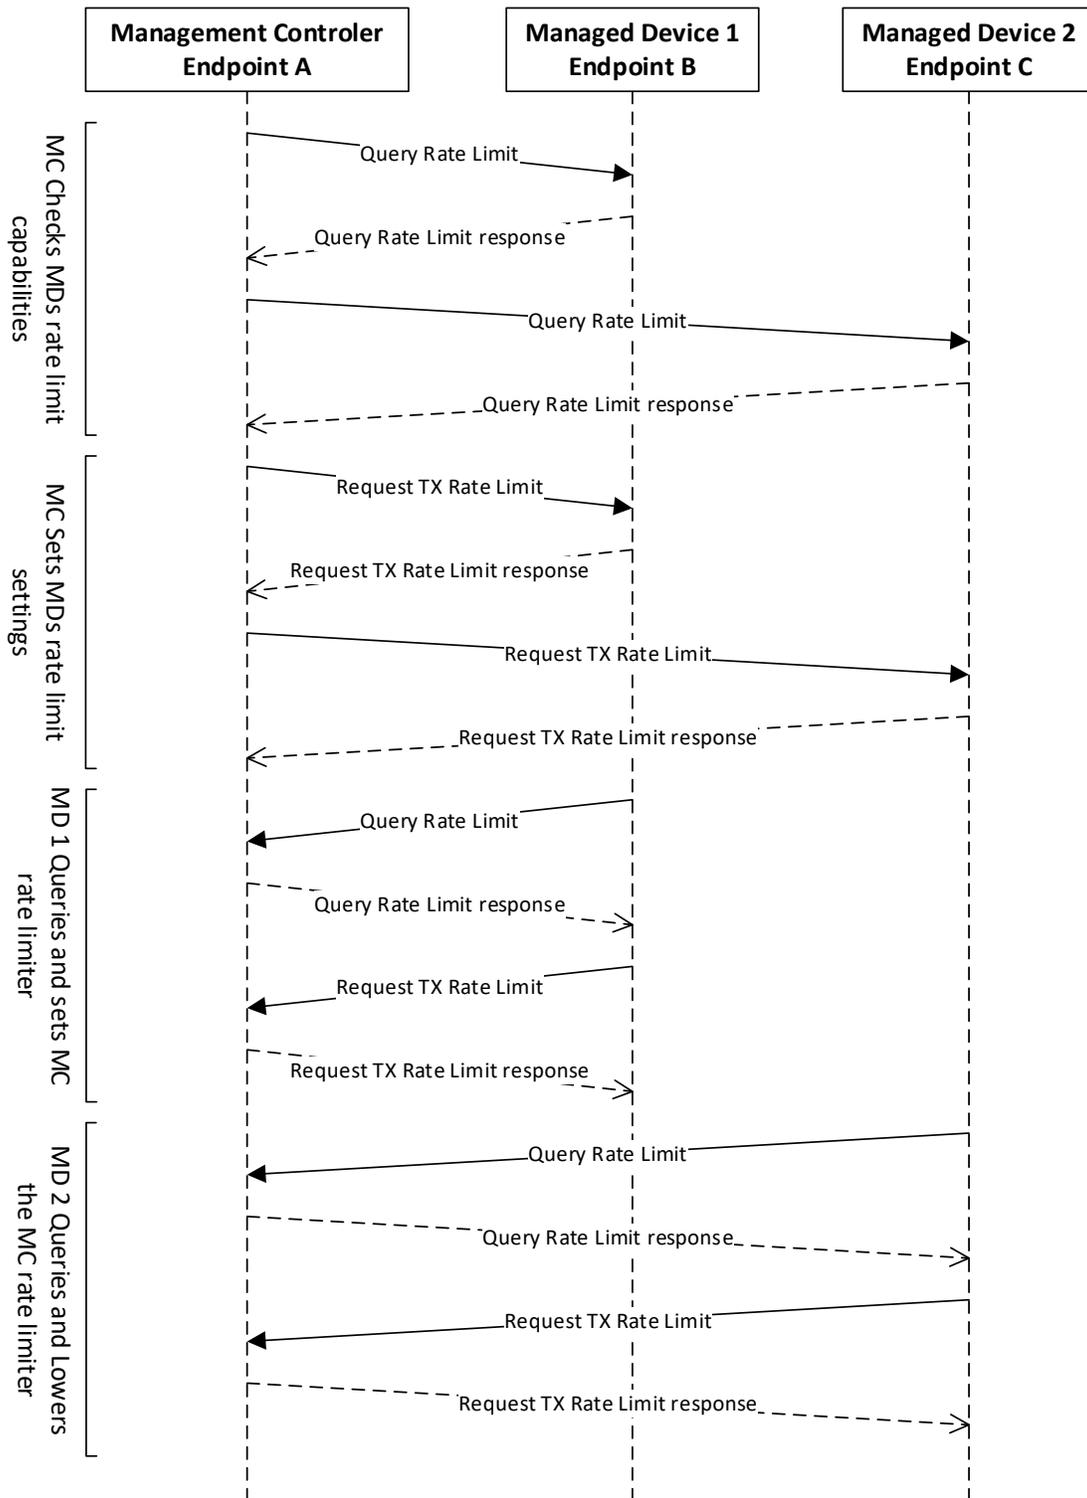
- 1749 • Publish its input processing rate and whether it can rate limit its output
- 1750 • Request its partner to rate limit its MCTP output traffic.

1751 Rate Limiting is negotiated between two endpoints and is configured on a per-EID basis such that devices
1752 having multiple EIDs should separately negotiate their Rate Limiting for each EID which supports Rate
1753 Limiting. If there are any MCTP Bridges in the path between the endpoints, the negotiated rate limit
1754 between the endpoints may not take bridge performance into account. The negotiation should take the
1755 speeds of the media for the path between the endpoints into account. Rate limiting is not specified for the
1756 bridging functionality within an MCTP Bridge (the functionality that routes MCTP packets between
1757 different ports on the bridge).

1758 Figure 19 presents an example of message exchanges for rate limiting. In this example, the management
1759 controller (MC) wants the managed devices (MD 1, MD 2) to send data at a limited rate. The MC first
1760 queries the MDs for their rate limiting capabilities using the Query Rate Limit command. Based on those
1761 capabilities, the MC requests the maximal transmit rate configuration for the MDs using the Request TX
1762 rate limit command. Conversely, the MDs may want to limit the rate that they receive data from the MC. In
1763 this case, it's the MDs that query the MC using the Query Rate Limit command, and, based on the
1764 response from the MC, requests configuration for the transmit rate from the MC using the Request TX
1765 rate limit command.

1766 Note that the figure does not show conditions such as handling the situation where one or more of the
1767 endpoints does not support Rate Limiting, nor does it show any algorithms that the endpoints may use to
1768 determine the best end-to-end value for Rate Limiting. Devices that negotiate Rate Limiting may wish to
1769 include algorithms or tests that would indicate there are intermediate devices in the path, such as
1770 Bridges, that would require transmit rates to be set to values that are lower than just what the receiving
1771 device needs. For example, the receiving device may detect that additional Rate Limiting is needed by
1772 noticing that there are packets missing in a multi-packet MCTP message transfer sequence.

1773



1774

1775

Figure 19 – Example rate limiting message exchanges

1776 **10.1.2 Restrictions on rate limiting**

1777 Message-based flow control may not utilize rate limiting. When rate limiting is active on a device which
1778 sends non-requested messages, then request/responses may also be affected by the rate limiting. Rate-
1779 limiting capable device may use rate limiting only to non-requested messages or to all messages. The
1780 transmit rate limiting operation-mode capability is reported by the device through "Transmit Rate Limiting
1781 operation capability" bit in Query rate limit command response.

1782 The use of rate limiting shall not supersede the timing requirements that are called out in other
1783 specifications, such as the transport binding specifications. Rate limiting shall include configuration
1784 options that allow meeting timing requirements under nominal operating conditions.

1785 **10.1.3 Rate definition**

1786 Let B be the Maximum supported burst size and R be the Maximum output rate limit in Packets Per
1787 Second (PPS), then the traffic shall be throttled such that in **any** time window $W = B/R$ (where $B \geq$
1788 1) there are no more than B packets.

1789 **10.1.4 Output rate limiting capabilities parameters**

1790 A transmitter that supports rate limiting shall expose its rate limiting capabilities using the Query Rate
1791 Limit command. For the definition of rate limiting, a baseline-transmission packet includes the baseline
1792 transmission unit as well as any medium-specific header/trailer and MCTP transport header. This
1793 includes:

- 1794 • **Maximum output rate limit:** The maximum rate in baseline transmission unit Packets/sec that
1795 the transmitting endpoint can be limited to when sending data to another endpoint.
- 1796 • **Minimum output rate limit:** The minimum rate in baseline transmission unit Packets/sec that
1797 the transmitting endpoint can be limited to when sending data to another endpoint. This value is
1798 also used to define the granularity of the configurable rate limit values.
- 1799 • **Maximum supported burst size:** The maximum number of consecutive baseline transmission
1800 unit Packets that the transmitter endpoint can send with minimal delay between MCTP packets.

1801 **10.1.5 Input processing capabilities parameters**

1802 A receiver can expose its input processing capabilities using the Query Rate Limit command. These
1803 parameters are informative only and should not be used to set the rate limiter of the partner. These
1804 parameters are intended to be used for visibility on the transmitter side, for performance analysis and
1805 monitoring purpose.

1806 The parameters exposed are:

- 1807 • **Maximum allowed receive data rate:** The maximum processing rate in baseline transmission
1808 unit packets/sec that the receiving endpoint can typically process incoming traffic. The data rate
1809 is measured using a time window. This rate is defined regardless of the content being received.
1810 Thus, devices which are limited in message processing shall report the maximum allowed
1811 receive data rate for minimal-size packets.
- 1812 • **Buffer Size:** this parameter defines the receive **buffer size** in bytes of the receiving endpoint.

1813 10.1.6 Defining and updating configuration parameters

1814 10.1.6.1 Rate limiting configuration parameters

1815 Rate limiting requirements are defined explicitly for each endpoint by means of two parameters, the
1816 maximum allowed data rate and the maximum continuous burst size. These are defined as follows:

- 1817 • **Maximum continuous burst size:** The maximum continuous burst size is defined in MCTP
1818 packets. Typically, this parameter reflects the receive buffer resources of the receiving endpoint.
- 1819 • **Maximum allowed data rate:** The maximum allowed data rate is defined in baseline
1820 transmission unit packets/sec. Typically, this defines the rate at which a receiving endpoint can
1821 process incoming messages. The data rate is measured using a time window as defined above.
1822 This rate is defined regardless of the content being received. Thus, devices which are limited in
1823 message processing shall request the maximum allowed transmit data rate with Burst Size of 1
1824 packet.

1825 If a device contains more than one MCTP endpoint (for example, a device that has an endpoint on
1826 SMBus/I2C and one on PCIe VDM) and supports setting rate limiting on these endpoints, then each rate-
1827 limiting configuration shall be independent and separately configurable. A device may include rate limiting
1828 capability for part or all of the endpoints.

1829 These parameters are used both by the receiver to request a specific traffic rate from the transmitter
1830 device and by the transmitter device to report the current rate-limiting values.

1831 When different settings are requested from different receiving endpoints, a transmitting endpoint that
1832 implements a single rate limiter shall use the smallest continuous burst size and the lowest data-rate that
1833 has been requested across the set of receiving endpoints. In a case of a single rate limiter, when traffic to
1834 multiple EIDs is active at the same time, the effective data rate to each of the receiving EIDs may be
1835 lower than the configured rate, as the aggregated data rates to all receiving EIDs will be the configured
1836 Rate Limiting settings.

1837 When a system is designed with devices supporting rate-limiting and devices which do not support rate-
1838 limiting, any device which supports rate limiting shall set its rate limiter to the negotiated rate-limiting
1839 settings. It is recommended that devices which do not support rate limiting are configured such that they
1840 will not cause buffer-overflow or data-processing rate overflow to their connected receiving endpoint. The
1841 implementation method of such a system is outside the scope of this specification.

1842 10.1.6.2 Updating rate-limiting parameters

1843 If an endpoint device needs to update the rate limiting settings of the other endpoint devices which are
1844 communicating with it and which are configured with rate limiting, it shall request the new settings in the
1845 sending devices using Request TX rate limit command. Once the response to Request TX rate limit
1846 command is received, the new rate limit is set according to the settings provided in the response. When
1847 the rate limiting settings is changed by an endpoint, the transmitting endpoint should notify the other
1848 receiving endpoint, sharing the same rate limiter, about the update using the Update rate limit command.

1849 11 Resiliency and recovery

1850 11.1 Overview

1851 Being a management fabric, MCTP bus and MCTP endpoints may require resetting an Endpoint, a
1852 physical bus or a firmware stack on a managed device. For this reason, certain reset methods are defined
1853 as optional methods to improve the resiliency of a managed system when a given reset type is required.

1854 In some cases, MCTP bridges are used, and these bridges abstract the buses to which they are
1855 connected. In these cases, there is a need to allow native recovery methods for these buses which are
1856 not directly accessible through the top-most MCTP bus owner.

1857 **11.2 MCTP reset types**

1858 MCTP defined multiple levels of optional resets.

1859 **11.2.1 MCTP application-level reset (L1)**

1860 Application-level reset is intended to reset a given upper-layer protocol such as PLDM, NC-SI and any
1861 other protocol that uses MCTP as the transport protocol.

1862 This reset can be used to reset a software stack and/or initialize runtime data structures which may
1863 require initialization to return to functional state.

1864 Application-level reset shall not affect the MCTP transport layer or other management protocols using
1865 MCTP on the same device.

1866 **11.2.2 MCTP transport reset (L2)**

1867 MCTP Transport Reset assumes that the Device implementation has enough capabilities to ensure that
1868 subsequent MCTP Transfers are handled as expected without need for further intervention. This includes,
1869 but is not limited to, flushing internal pipelines, clearing timeout flags and resetting counters etc. when an
1870 endpoint error condition is cleared.

1871 **11.2.3 Physical bus-level reset (L3)**

1872 Physical bus-level reset is expected to be detected by a dedicated HW of the device's bus-interface
1873 hardware (e.g., USB PHY for USB 2.0 High Speed). Bus-specific reset detection mechanism shall be
1874 isolated from the endpoint physical interface logic, ensuring reliable reset event recognition, regardless of
1875 the current bus-level communication state.

1876 This reset type is bus-specific and may be implemented only on specific bus types.

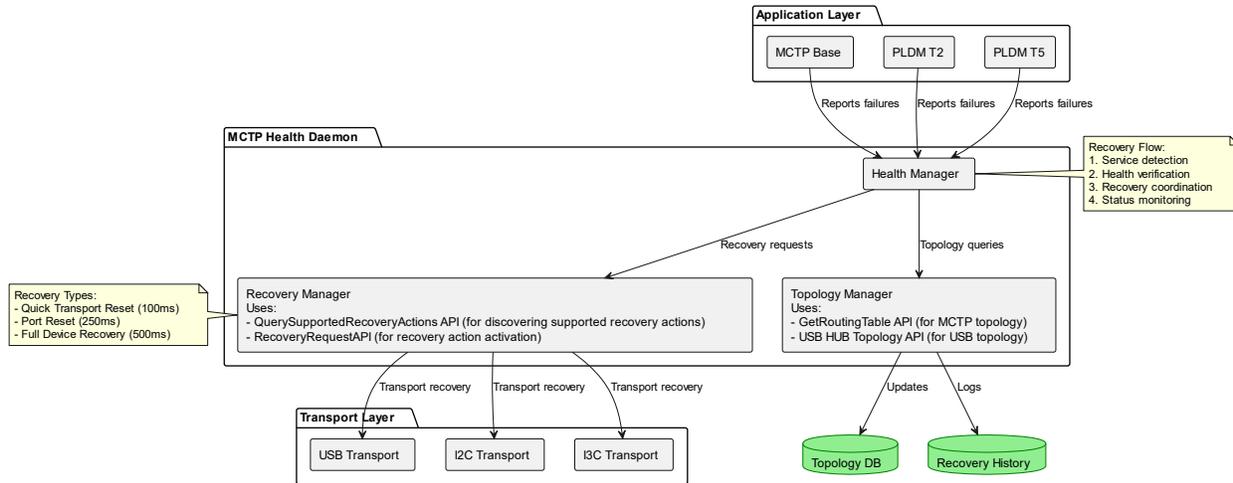
1877 **11.2.4 Device manageability subsystem reset (L4)**

1878 This reset type includes reset of manageability subsystems of the device identified by the MCTP Endpoint
1879 ID. It should be implemented either by a subsystem soft-reset or as a hardware reset. It is expected that
1880 reset of the manageability subsystem will not impact on any other runtime functionality of the Device.
1881 When supported, this reset type shall recover the manageability device subsystem.

1882 **11.3 MCTP System Resiliency**

1883 An MCTP system which supports MCTP resiliency, has to learn the supported resiliency methods of
1884 every endpoint and every Bridge. The resiliency-support query for endpoints shall be performed following
1885 every enumeration of the endpoints. MCTP bridge's resiliency support shall include both the bridge
1886 endpoints and the downstream MCTP ports resiliency methods.

1887 The following diagram illustrates an MCTP system with resiliency support.



1888

1889

Figure 20 – MCTP system with resiliency support

1890 The system in Figure 20 uses multiple upper-level applications such as PLDM T2 and PLDM T5. The
 1891 health daemon entity should be included in the top-level MCTP bus owner. The Health Manager is an
 1892 example of a SW entity which is used to monitor the system’s health and initiate the required supported
 1893 resiliency actions if required.

1894 The Recovery manager is responsible for activating the required action as applicable. Note that the
 1895 required action and the supported actions may differ between different MCTP physical bus-types. Bus-
 1896 specific resiliency actions should be defined in the corresponding binding specifications for every bus
 1897 type.

1898 **12 MCTP control protocol**

1899 **12.1.1 Overview**

1900 MCTP control messages are used for the setup and initialization of MCTP communications within an
 1901 MCTP network. This clause defines the protocol and formatting used for MCTP control messages over
 1902 MCTP.

1903 **12.2 Terminology**

1904 The terms shown in Table 8 are used when describing the MCTP control protocol.

1905

Table 8 – MCTP control protocol terminology

Term	Description
Requester	The term “requester” is used to refer to the endpoint that originates an MCTP control Request message.
Responder	The term “responder” is used to refer to the endpoint that originates an MCTP control response message (that is, an endpoint that returns the response to an MCTP control Request message).
Originator or Source	The term “originator” or “source” is used to refer to the endpoint that originates any MCTP control message: Request, Response, or Datagram.
Target or Destination	The term “target” or “destination” is used to refer to the endpoint that is the intended recipient of any MCTP control message: Request, Response, or Datagram.

Term	Description
Asynchronous Notification	The term “asynchronous notification” is used to refer to the condition when an MCTP endpoint issues an un-requested Datagram to another MCTP endpoint.
Broadcast	The term “broadcast” is used when an MCTP control Datagram is sent out onto the bus using the broadcast EID.

1906 12.3 Control message classes

1907 12.3.1 General

1908 The different types of messages shown in Table 9 are used under the MCTP control message type.

1909 **Table 9 – MCTP control message types**

Type	Description
Request	This class of control message requests that an endpoint performs a specific MCTP control operation. All MCTP control Request messages are acknowledged with a corresponding Response message. (Within this specification, the term “command” and “request” are used interchangeably as shorthand to refer to MCTP control Request messages.)
Response	This class of MCTP control message is sent in response to an MCTP control Request message. The message includes a “Completion Code” field that indicates whether the response completed normally. The response can also return additional data dependent on the particular MCTP control Request that was issued. An MCTP control Response message shall use the destination EID and physical address that were used as the source EID and source physical address of the corresponding MCTP control Request message.
Datagram	Datagrams are “unacknowledged” messages (that is, Datagrams do not have corresponding Response messages). This class of MCTP control message is used to transfer messages when an MCTP control Response message is neither required nor desirable.
Broadcast Request	A broadcast message is a special type of Request that is targeted to all endpoints on a given bus. All endpoints that receive the message are expected to interpret the Request.
Broadcast Datagram	A Datagram that is broadcast to all endpoints on the bus. Broadcast Datagrams are “unacknowledged” messages (that is, broadcast Datagrams do not have corresponding Response messages).

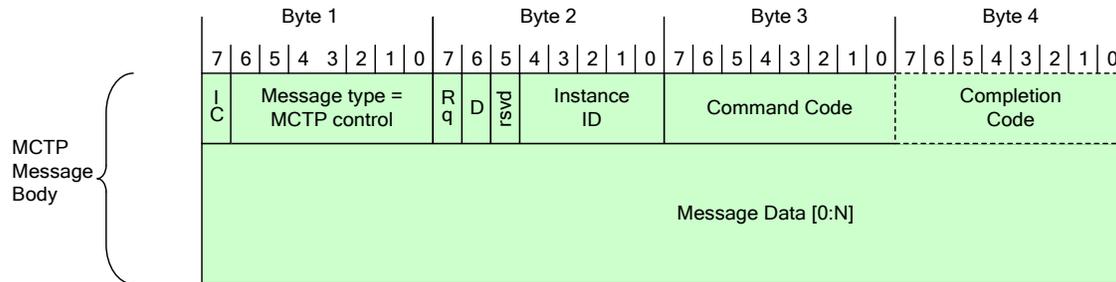
1910 12.4 MCTP control message format

1911 12.4.1 Overview

1912 MCTP control messages use the MCTP control message type (see Table 3). Any message sent with this
 1913 message type will correspond to the definitions set forth in this clause. The basic format of an MCTP
 1914 control message is shown in Figure 21. Note that the byte offsets shown in Figure 21 are relative to the
 1915 start of the MCTP message body rather than the start of the physical packet.

1916 12.4.2 Use of Message Integrity Check

1917 MCTP control messages do not use a Message Integrity Check field. Therefore, the IC bit in MCTP
 1918 control messages shall always be 0b.



1919

1920

Figure 21 – MCTP control message format

1921 **12.5 MCTP control message fields**

1922 Table 10 lists the common fields for MCTP control messages.

1923

Table 10 – MCTP control message fields

Field Name	Description
IC*	Message Integrity Check bit = 0b. MCTP control messages do not include an overall Message Integrity check field.
Message Type*	MCTP control = 0x00 (000_0000b). This field identifies the MCTP message as being an MCTP control message.
Rq bit	Request bit. This bit is used to help differentiate between MCTP control Request messages and other message classes. Refer to 12.7.
D-bit	Datagram bit. This bit is used to indicate whether the Instance ID field is being used for tracking and matching requests and responses or is just being used to identify a retransmitted message. Refer to 12.7.
Instance ID	The Instance ID field is used to identify new instances of an MCTP control Request or Datagram to differentiate new requests or datagrams that are sent to a given message terminus from retried messages that are sent to the same message terminus. The Instance ID field is also used to match up a particular instance of an MCTP Response message with the corresponding instance of an MCTP Request message.
Command Code	For Request messages, this field is a command code indicating the type of MCTP operation the packet is requesting. Command code values are defined in Table 12. The format and definition of request and response parameters for the commands is given in Clause 13. The Command Code that is sent in a Request shall be returned in the corresponding Response.

Field Name	Description
Completion Code	This field is only present in Response messages. This field contains a value that indicates whether the response completed normally. If the command did not complete normally, the value can provide additional information regarding the error condition. The values for completion codes are specified in Table 13.
Message Data	Zero or more bytes of parameter data that is specific to the particular Command Code and whether the message is a Request or Datagram, or a Response.
* These fields are MCTP base protocol fields.	

1924 **12.6 MCTP control message transmission unit size**

1925 All MCTP control messages are required to have a packet payload that is no larger than the baseline
 1926 transmission unit size of 64 bytes.

1927 MCTP control messages are carried in a single MCTP packet. Multiple messages are used if an operation
 1928 requires more data to be transferred than can be carried in a single message.

1929 **12.7 Tag Owner (TO), Request (Rq), and Datagram (D) bit usage**

1930 For MCTP control messages, the Rq bit shall be set to 1b if the message is a “command” or Request
 1931 message and 0b if the message is a Response message. For Datagram and Broadcast messages, the
 1932 Rq bit shall always be set to 1b. MCTP Control messages that have unexpected or incorrect flag bit
 1933 values shall be silently discarded by the receiver of the message.

1934 For the present specification, Requests and Datagrams are only issued from tag owners (TO bit = 1b).
 1935 Provision has been left for the definition of possible future Datagrams that are not issued from tag owners
 1936 (see Table 11).

1937 **Table 11 – Tag Owner (TO), Request (Rq) and Datagram (D) bit usage**

MCTP Control Message Class	Destination EID Value	Tag Owner (TO) bit	Request (Rq) bit	Datagram (D) bit
Command/Request Responses are expected and tracked by Instance ID at the requester.	Target EID	1b	1b	0b
Response	Target EID	0b	0b	0b
Broadcast Request Responses are expected and tracked by Instance ID at the requester.	Broadcast EID	1b	1b	0b
Datagram Unacknowledged Request – Responses are neither expected nor tracked by Instance ID at the requester. Duplicate packets are handled the same as retried Command/Request packets.	Target EID	1b	1b	1b
Broadcast Datagram (unacknowledged control command that is broadcast.)	Broadcast EID	1b	1b	1b
Reserved for future definition	all other			

1938 **12.8 Concurrent command processing**

1939 **12.8.1 Overview**

1940 This clause describes the specifications and requirements for handling concurrent overlapping MCTP
1941 control requests by endpoints.

1942 **12.8.2 Requirements for responders**

1943 An endpoint is not required to process more than one request at a time (that is, it can be “single threaded”
1944 and does not have to accept and act on new requests until it has finished responding to any previous
1945 request).

1946 A responder that is not ready to accept a new request can either silently discard the request, or it can
1947 respond with an `ERROR_NOT_READY` message completion code.

1948 A responder that can accept and process more than one request at a time is not required to return
1949 responses in the order that the requests were received.

1950 **12.8.3 Requirements for Requesters**

1951 An endpoint that issues MCTP control Requests to another endpoint shall wait until it gets the response
1952 to the particular request, or times out waiting for the response, before issuing a new request, Datagram,
1953 or Broadcast Datagram.

1954 An endpoint that issues MCTP control Requests is allowed to have multiple requests outstanding
1955 simultaneously to *different* responder endpoints.

1956 An endpoint that issues MCTP control Requests should be prepared to handle responses that may not
1957 match the request (that is, it should not automatically assume that a response that it receives is for a
1958 particular request). It should check to see that the command code and source EID values in the response
1959 match up with a corresponding outstanding command before acting on any parameters returned in the
1960 response.

1961 **12.8.4 Additional requirements for bridges**

1962 The packets that are routed *through* a bridge’s routing functionality are not interpreted by the bridge and
1963 therefore are not considered to constitute concurrent requests.

1964 A bridge shall support at least one outstanding MCTP control request for each bus connection (port)
1965 through which MCTP control messages can be used to access the bridge’s configuration and control
1966 functionality.

1967 Bridges shall retain temporal ordering of packets forwarded from one message terminus to another.

1968 **13 MCTP control messages**

1969 **13.1 Overview**

1970 This clause contains detailed descriptions for each MCTP control message. The byte offsets for the
1971 Request and Response parameter information given in the tables for the commands indicates the byte
1972 offset for the message data starting with the byte following the Command field.

1973 **13.2 MCTP control message command codes**

1974 Table 12 lists the MCTP control messages and their corresponding command code values. The
 1975 commands and their associated parameters are specified later in this clause. For bridges, the
 1976 requirements apply equally to all endpoints within the bridge device that are used to configure and control
 1977 the bridges routing functionality.

1978 **Table 12 – MCTP control command numbers**

Command Code	Command Name	General Description	OMC		Clause
			E	B	
0x00	Reserved	Reserved	–	–	–
0x01	Set Endpoint ID	Assigns an EID to the endpoint at the given physical address	Ma Ng	Ca ¹ Mg	13.4
0x02	Get Endpoint ID	Returns the EID presently assigned to an endpoint. Also returns information about what type the endpoint is and its level of use of static EIDs.	Ma Og	Ma Og	13.5
0x03	Get Endpoint UUID	Retrieves a per-device unique UUID associated with the endpoint	Ca ² Og ⁹	Ca ² Og	13.6
0x04	Get MCTP Version Support	Lists which versions of the MCTP control protocol are supported on an endpoint	Ma Og	Ma Og ⁵	13.7
0x05	Get Message Type Support	Lists the message types that an endpoint supports	Ma Og	Ma Og	13.8
0x06	Get Vendor Defined Message Support	Used to discover an MCTP endpoint's vendor-specific MCTP extensions and capabilities	Oa Og	Oa Og	13.9
0x07	Resolve Endpoint ID	Used to get the physical address associated with a given EID	Na Og	Ma Og	13.10
0x08	Allocate Endpoint IDs	Used by the bus owner to allocate a pool of EIDs to an MCTP bridge	Na Ng	Ma ⁶ Mg ⁶	13.11
0x09	Routing Information Update	Used by the bus owner to extend or update the routing information that is maintained by an MCTP bridge	Oa ⁸ Og ⁸	Ma ⁴ Mg ⁴	13.12
0x0A	Get Routing Table Entries	Used to request an MCTP bridge to return data corresponding to its present routing table entries	Na Og	Ma Og	13.13
0x0B	Prepare for Endpoint Discovery	Used to direct endpoints to clear their "discovered" flags to enable them to respond to the Endpoint Discovery command	Ca ³ Ng	Ca ³ Cg ³	13.14
0x0C	Endpoint Discovery	Used to discover MCTP-capable devices on a bus, provided that another discovery mechanism is not defined for the particular physical medium	Ca ³ Cg ³	Ca ³ Cg ³	13.15
0x0D	Discovery Notify	Used to notify the bus owner that an MCTP device has become available on the bus	Na Cg ³	Ca ³ Cg ³	13.16
0x0E	Get Network ID	Used to get the MCTP network ID	Ca ⁷	Ca ⁷	13.17
0x0F	Query Hop	Used to discover what bridges, if any, are in the path to a given target endpoint and what transmission unit sizes the bridges will pass for a given message type when routing to the target endpoint	Na Og	Ma Og	13.18
0x10	Resolve UUID	Used by endpoints to find another endpoint matching an endpoint that uses a specific UUID.	Na Og	Oa Og	13.19

Command Code	Command Name	General Description	OMC		Clause
			E	B	
0x11	Query rate limit	Used to discover the data rate limit settings of the given target for incoming messages.	Oa Og	Oa Og	13.20
0x12	Request TX rate limit	Used to request the allowed transmit data rate limit for the given endpoint for outgoing messages.	Oa Og	Oa Og	13.21
0x13	Update rate limit	Used to update the receiving side on change to the transmit data rate which was not requested by the receiver	Oa Og	Oa Og	13.22
0x14	Query Supported Interfaces	Used to discover the existing device MCTP interfaces.	Oa Og	Oa Og	13.23
0x15	Query Endpoint Advanced Capabilities	Used to query the advanced capabilities of a DSP0236 1.4.0 (or later) device	Ma Og	Oa Og	13.24
0x1A	Request EID pool	Used by an MCTP bridge to request allocation of additional EIDs to a requester bridge	Na Ng	Oa Og	13.25
0x1B	Release EID pool	Used by an MCTP bus owner to request a lower-level bridge to free unassigned EIDs	Na Ng	Oa Og	13.26
0x1C	Query Supported Recovery Actions	Used by MCTP bus owners query supported recovery actions	Oa Og	Oa Og	0
0x1D	Recovery Request	Used by MCTP bus owners to try to recover connection to an MCTP endpoint	Oa Og	Oa Og	13.28
0xF0 - 0xFF	Transport Specific	This range of control command numbers is reserved for definition by individual MCTP Transport binding specifications. Transport specific commands are intended to be used as needed for setup and configuration of MCTP on a given media. A particular transport specific command number many have different definitions depending on the binding specification. Transport specific commands shall only be addressed to endpoints on the same medium. A bridge is allowed to block transport specific commands from being bridged to different media. The general format of Transport specific messages is specified in clause 13.18.	-	-	13.24

Key for OMC (optional / mandatory / conditional) column:

E = non-bridge, non-bus owner endpoint (simple endpoint)
 B = bridge / bus-owner endpoint
 Ma = mandatory (required) to accept. The request shall be accepted by the endpoint and a response generated per the following command descriptions.
 Mg = mandatory to generate. The endpoint shall generate this request as part of its responsibilities for MCTP operation.
 Oa = optional to accept
 Og = optional to generate
 Ca = conditional to accept (see notes)
 Cg = conditional to generate (see notes)
 Na = not applicable to accept. This command is not applicable to the device type and shall not be accepted
 Ng = not applicable to generate. This command is used for MCTP configuration and initialization and should not be generated.

1. The topmost bus owner is not required to support the Set Endpoint ID command.
2. Hot-plug and add-in devices, and non-bridge devices that connect to multiple busses, are required to support the Get Endpoint UUID command. See 8.18.10 and 8.18.11 for more info.

Command Code	Command Name	General Description	OMC		Clause
			E	B	
		<ol style="list-style-type: none"> 3. Mandatory on a per-bus basis to support endpoint discovery if required by the physical transport binding used for the particular bus type. Refer to the appropriate MCTP physical transport binding specification. 4. The topmost bus owner is not required to accept this command. The command is required to be generated when downstream bridges require dynamic routing information from bus owners that they are connected to. Some implementations may be configured where all routing information has been statically configured into the bridge and no dynamically provided information is required. In this case, it is not required to support the command while the endpoints are configured in that manner. 5. Bridges should use this command to verify that they are initializing devices that are compatible with their MCTP control protocol version. 6. The endpoint is required to accept this command if it indicated support for a dynamic EID pool. The command shall be generated by the endpoint if the configuration requires the endpoint to support allocating EID pools to downstream bridges. 7. See Clause 9 for more information regarding MCTP Network IDs and 12.17 regarding the implementation requirements of this command. 8. While it is optional for an endpoint to receive a routing information update, the MCTP Base specification does not specify a bridge or bus owner function that sends such updates to particular endpoints. 9. While it is optional for an endpoint to support this command, support of this command is mandatory both to generate and to accept for devices supporting rate limiting. 			

1979 13.3 MCTP control message completion codes

1980 The command/result code field is used to return management operation results for response messages. If
 1981 a `SUCCESS` completion code is returned then the specified response parameters (if any) shall also be
 1982 returned in the response. If an error completion code (not `SUCCESS`) is returned by the responder, unless
 1983 otherwise specified, the responder shall not return any additional parametric data and the requester shall
 1984 ignore any additional parameter data provided in the response (if any). See Table 13 for the completion
 1985 codes.

1986 **Table 13 – MCTP control message completion codes**

Value	Name	Description
0x00	<code>SUCCESS</code>	The Request was accepted and completed normally.
0x01	<code>ERROR</code>	This is a generic failure message. (It should not be used when a more specific result code applies.)
0x02	<code>ERROR_INVALID_DATA</code>	The packet payload contained invalid data or an illegal parameter value.
0x03	<code>ERROR_INVALID_LENGTH</code>	The message length was invalid. (The Message body was larger or smaller than expected for the particular request.)
0x04	<code>ERROR_NOT_READY</code>	The Receiver is in a transient state where it is not ready to receive the corresponding message.
0x05	<code>ERROR_UNSUPPORTED_CMD</code>	The command code field of the received message is unspecified or not supported on this endpoint. This completion code shall be returned for any unsupported command values received in MCTP control Request messages.
0x80–0xFF	<code>COMMAND_SPECIFIC</code>	This range of completion code values is reserved for values that are specific to a particular MCTP control message. The particular values (if any) and their definition is provided in the specification for the particular command.
All other	Reserved	Reserved

1987 **13.4 Set Endpoint ID**

1988 The Set Endpoint ID command assigns an EID to an endpoint and sets its Discovered Flag for the
 1989 endpoint on the physical bus from which this command was received. This command should only be
 1990 issued by a bus owner to assign an EID to an endpoint at a particular physical address. Since it is
 1991 assumed the Endpoint does not already have an EID assigned to it, or because the EID is unknown, the
 1992 destination EID in the message will typically be set to the special null destination EID value.

1993 The Set Endpoint ID command is also used to provide the Physical Address and EID of the Bus Owner to
 1994 an Endpoint. An Endpoint that needs to communicate with the Bus Owner may capture the physical
 1995 address and EID that was used to deliver the Set Endpoint ID message.

1996 Note: Endpoints that are not the Bus Owner should not issue the Set Endpoint ID command because it can
 1997 cause the receiver of the message to capture incorrect information for the Bus Owner's address.

1998 An MCTP bridge may elect to have a single EID for its functionality, rather than using an EID for each port
 1999 (bus connection) that is connected to a different bus owner. See 9.1.3 for more information. In this case,
 2000 the bridge will accept its EID assignment from the "first" bus to deliver the Set Endpoint ID request to the
 2001 bridge.

2002 It is recognized that different internal processing delays within a bridge can cause the temporal ordering
 2003 of requests to be switched if overlapping requests are received over more than one bus. Therefore, which
 2004 request is accepted by an implementation is not necessarily tied to the request that is first received at the
 2005 bridge, but instead will be based on which request is the first to be processed by the bridge.

2006 If an EID has already been assigned and the Set Endpoint ID command is issued to set a new EID, the
 2007 command shall return a `SUCCESSFUL` completion code, and the response shall use the same EID value
 2008 that was used as the destination EID in the Set Endpoint ID command.

2009 If an EID has already been assigned and the Set Endpoint ID command is issued from a different bus
 2010 without forcing an EID assignment, the command shall return a `SUCCESSFUL` completion code, but the
 2011 response parameters shall return an EID assignment status of "EID rejected".

2012 The Set Endpoint ID command functions in the same manner regardless of whether the endpoint uses a
 2013 static EID. The only difference is that if an endpoint has a static EID, it uses that EID as its initial "default"
 2014 EID value. The endpoint does not treat this initial EID as if it were assigned to it by a different bus owner.
 2015 That is, the endpoint shall accept the EID assignment from the first bus that the command is received
 2016 from, and shall track that bus as the originating bus for the EID for subsequent instances of Set Endpoint
 2017 ID command. See 8.18.3 for more information. The request and response parameters are specified in
 2018 Table 14.

2019 **Table 14 – Set Endpoint ID message**

	Byte	Description
Request data	1	Operation [7:2] – reserved [1:0] – Operation: 00b Set EID. Submit an EID for assignment. The given EID will be accepted conditional upon which bus the device received the EID from (see preceding text). A device where the endpoint is only reached through one bus shall always accept this operation (provided the EID value is legal). 01b Force EID. Force EID assignment. The given EID will be accepted regardless of whether the EID was already assigned through another bus. Note that if the endpoint is forcing, the EID assignment changes which bus is being

	Byte	Description
		<p>tracked as the originator of the Set Endpoint ID command. A device where the endpoint is only reached through one bus shall always accept this operation (provided the EID value is legal), in which case the Set EID and Force EID operations are equivalent.</p> <p>10b Reset EID (optional). This option only applies to endpoints that support static EIDs. If static EIDs are supported, the endpoint shall restore the EID to the statically configured EID value. The EID value in byte 2 shall be ignored. An <code>ERROR_INVALID_DATA</code> completion code shall be returned if this operation is not supported.</p> <p>11b Set Discovered Flag. Set Discovered flag to the "discovered" state only. Do not change present EID setting. The EID value in byte 2 shall be ignored. Note that Discovered flag is only used for some physical transport bindings. An <code>ERROR_INVALID_DATA</code> completion code shall be returned if this operation is selected and the particular transport binding does not support a Discovered flag.</p>
	2	<p>Endpoint ID. <code>0xFF</code>, <code>0x00</code> = illegal. Endpoints are not allowed to be assigned the broadcast or null EIDs. It is recommended that the endpoint return an <code>ERROR_INVALID_DATA</code> completion code if it receives either of these values.</p>
Response data	1	Completion code
	2	<p>[7:6] – reserved [5:4] – EID assignment status: 00b = EID assignment accepted. 01b = EID assignment rejected. EID has already been assigned by another bus owner and assignment was not forced. 10b = reserved. 11b = reserved. [3:2] – reserved. [1:0] – Endpoint ID allocation status (see 13.11 for additional information): 00b = Device does not use an EID pool. 01b = Endpoint requires EID pool allocation. 10b = Endpoint uses an EID pool and has already received an allocation for that pool. 11b = reserved</p>
	3	<p>EID Setting. If the EID setting was accepted, this value will match the EID passed in the request. Otherwise, this value returns the present EID setting.</p>
	4	<p>EID Pool Size. This is the size of the dynamic EID pool that the bridge can use to assign EIDs or EID pools to other endpoints or bridges. It does not include the count of any additional static EIDs that the bridge may maintain. See 8.18.3 for more information. Note that a bridge always returns its pool size regardless of whether it has already received an allocation. <code>0x00</code> = no dynamic EID pool.</p>

2020 **13.5 Get Endpoint ID**

2021 The Get Endpoint ID command returns the EID for an endpoint. This command is typically issued only by
 2022 a bus owner to retrieve the EID that was assigned to a particular physical address. Thus, the destination
 2023 EID in the message will typically be set to the special Physical Addressing Only EID value. The request
 2024 and response parameters are specified in Table 15.

2025 **Table 15 – Get Endpoint ID message**

	Byte	Description
Request data	–	–
Response data	1	Completion Code.
	2	Endpoint ID. 0x00 = EID not yet assigned.
	3	<p>Endpoint Type. [7:6] = reserved [5:4] = Endpoint Type: 00b = simple endpoint 01b = bus owner/bridge 10b = reserved 11b = reserved</p> <p>[3:2] = reserved [1:0] = Endpoint ID Type: 00b = dynamic EID. The endpoint uses a dynamic EID only. 01b = static EID supported.</p> <p>The endpoint was configured with a static EID. The EID returned by this command reflects the present setting and may or may not match the static EID value.</p> <p>The following two status return values are optional. If provided, they shall be supported as a pair in place of the static EID support status return. It is recommended that this be implemented if the Reset EID option in the Set Endpoint ID command is supported.</p> <p>10b = static EID supported. Present EID matches static EID. The endpoint has been configured with a static EID. The present value is the same as the static value.</p> <p>11b = static EID supported. Present EID does not match static EID. Endpoint has been configured with a static EID. The present value is different than the static value.</p> <p>See 8.18.3 for more information.</p>
4	<p>Medium-Specific Information.</p> <p>This byte can hold additional information about optional configuration of the endpoint on the given medium, such as whether certain types of timing or arbitration are supported. This should only be used to report static information.</p> <p>This byte shall be returned as 0x00 unless otherwise specified by the transport binding.</p>	

2026 **13.6 Get Endpoint UUID**

2027 The Get Endpoint UUID command returns a universally unique identifier (UUID), also referred to as a
 2028 globally unique ID (GUID), for the management controller or management device. The command can be
 2029 used to correlate a device with one or more EIDs. The format of the ID follows the byte (octet) format
 2030 specified in [RFC4122](#). [RFC4122](#) specifies four different versions of UUID formats and generation
 2031 algorithms suitable for use for a device UUID in IPMI. These are version 1 (0001b) “time based”, and
 2032 three “name-based” versions: version 3 (0011b) “MD5 hash”, version 4 (0100b) “Pseudo-random”, and
 2033 version 5 “SHA1 hash”. The version 1 format is recommended. However, versions 3, 4, or 5 formats are
 2034 also allowed. A device UUID should never change over the lifetime of the device. The request and
 2035 response parameters are specified in Table 16.

2036 See 8.18.10 and 8.18.11 for additional requirements on the use of the Get Endpoint UUID command.

2037 **Table 16 – Get Endpoint UUID message format**

	Byte	Description
Request data	–	–
Response data	1	Completion Code
	2:17	UUID bytes 1:16, respectively (see Table 17)

2038 The individual fields within the UUID are stored most-significant byte (MSB) first per the convention
 2039 described in [RFC4122](#). See Table 17 for an example format.

2040 **Table 17 – Example UUID format**

Field	UUID Byte	MSB
time low	1	MSB
	2	
	3	
	4	
time mid	5	MSB
	6	
time high and version	7	MSB
	8	
clock seq and reserved	9	MSB
	10	
node	11	MSB
	12	
	13	
	14	
	15	
	16	

2041 **13.7 Get MCTP version support**

2042 **13.7.1 Overview**

2043 This command can be used to retrieve the MCTP base specification versions that the endpoint supports,
 2044 and also the message type specification versions supported for each message type. The format of the
 2045 request and response parameters for this message is given in Table 18.

2046 More than one version number can be returned for a given message type by the Get MCTP Version
 2047 Support command. This enables the command to be used for reporting different levels of compatibility
 2048 and backward compatibility with different specification versions. The individual specifications for the given
 2049 message type define the requirements for which versions number values should be used for that
 2050 message type. Those documents define which earlier version numbers, if any, shall also be listed.

2051 The command returns a completion code that indicates whether the message type number passed in the
 2052 request is supported or not. This enables the command to also be used to query the endpoint for whether
 2053 it supports a given message type.

2054 NOTE Version numbers are listed from oldest to newest. Versioning commands and version formats for vendor-
 2055 defined message types, 0x7E and 0x7F, are vendor-specific and considered outside the scope of this specification.

2056 **Table 18 – Get MCTP version support message**

	Byte	Description
Request data	1	Message Type Number The Message Type Number to retrieve version information for: 0xFF = return MCTP base specification version information. 0x7E, 0x7F = unspecified. Support of this command for vendor-defined message types is vendor implementation-specific and considered outside the scope of this specification. 0x00 = return MCTP control protocol message version information. 0x01 = return version of DSP0241 0x02,0x03 = return version of DSP0261 Other = return version information for a given message type. See MCTP ID for message type numbers. When a Message Type Number references a binding spec, the reported version is of the binding spec and not of the associated base spec.
Response data	1	Completion Code 0x80 = message type number not supported
	2	Version Number Entry count One-based count of 32-bit version numbers being returned in this response. Numerically lower version numbers are returned first.
	3:6	Version Number entry 1: The following descriptions are informational. Refer to DSP4004 for the normative definition of version numbering of DMTF specifications. [31:24] = major version number. This field is used to identify a version of the specification that includes changes that make it incompatible with one or more functions that were defined in versions of the specification that have an older (smaller) major version number. [23:16] = minor version number. This field is used to identify functional additions to the specification that are backward compatible with

	Byte	Description
		<p>older (smaller) minor version numbers that share the same major version number.</p> <p>[15:8] = update version number. This field is used for editorial updates to the specification that do not define new functionality nor change existing functionality over the given major.minor release. This field is informational and should be ignored when checking versions for interoperability.</p> <p>[7:0] = “alpha” byte. This value is used for pre-release (work-in-progress) versions of the specification. Pre-release versions of the specification are backward compatible with specification versions that have an older (smaller) minor version numbers that share the same major version number. However, since the alpha value represents a version of the specification that is presently under development, versions that share the same major and minor version numbers, but have different ‘alpha’ versions may not be fully interoperable.</p> <p>The encoding of the version number and alpha fields is provided in 13.7.2.</p>
	(7:X)	<p>Version Number Entries 2 through N.</p> <p>Additional 32-bit major/minor version numbers, if any.</p> <p>This field is only included when there are 2 or more Version Number entries.</p>

2057 13.7.2 Version field encoding

2058 The version field is comprised of four bytes referred to as the “major”, “minor”, “update”, and “alpha”
 2059 bytes. These bytes shall be encoded as follows:

2060 The “major”, “minor”, and “update” bytes are BCD-encoded, and each byte holds two BCD digits. The
 2061 “alpha” byte holds an optional alphanumeric character extension that is encoded using one of the
 2062 alphabetic characters [a-z, A-Z] from the US-ASCII ([RFC20](#)) Character Set. The semantics of these fields
 2063 follows that specified in [DSP4004](#).

2064 The value 0x00 in the alpha field means that the alpha field is not used. Software or utilities that display
 2065 the version number should not display any characters for this field.

2066 The value 0xF in the most-significant nibble of a BCD-encoded value indicates that the most-significant
 2067 nibble should be ignored and the overall field treated as a single-digit value. Software or utilities that
 2068 display the number should only display a single digit and should not put in a leading “0” when displaying
 2069 the number.

2070 A value of 0xFF in the “update” field indicates that the field to be ignored. Software or utilities that display
 2071 the version number should not display any characters for the field. 0xFF is not allowed as a value for the
 2072 “major” or “minor” fields.

2073 EXAMPLES:

2074 Version 1.1.0 → 0xF1F1F000

2075 Version 3.1 → 0xF3F1FF00

2076 Version 1.0a → 0xF1F0FF61

2077 Version 3.7.10a → 0xF3F71061

2078 Version 10.11.7 → 0x1011F700

2079 13.7.3 MCTP base specification version number

2080 MCTP implementations that follow this particular specification shall return the following version
2081 information in the response to the Get MCTP Version Support message when the Message Type
2082 parameter in the request is set to 0xFF (return MCTP base specification version information).

2083 The Version Number Entry 1 field shall be used to indicate backward compatibility with Version 1.0 of the
2084 base specification as:

2085 1.0 [Major version 1, minor version 0, any update version, no alpha]

2086 This is reported using the encoding as: 0xF1F0FF00

2087 The Version Number Entry 2 field shall be used to indicate backward compatibility with Version 1.1 of the
2088 base specification as:

2089 1.1 [Major version 1, minor version 1, any update version, no alpha]

2090 This is reported using the encoding as: 0xF1F1FF00

2091 The Version Number Entry 3 field shall be used to indicate backward compatibility with Version 1.2 of the
2092 base specification as:

2093 1.2 [Major version 1, minor version 2, any update version, no alpha]

2094 This is reported using the encoding as: 0xF1F2FF00

2095 The Version Number Entry 4 field shall be used to indicate backward compatibility with Version 1.3 of the
2096 base specification as:

2097 1.3 [Major version 1, minor version 3, any update version, no alpha]

2098 This is reported using the encoding as: 0xF1F3FF00

2099 The version of the MCTP base specification for this specification shall be reported in Version Number
2100 Entry 5 as:

2101 **1.4.0** [Major version 1, minor version 4, update version 0, no alpha]

2102 This is reported using the encoding as: 0xF1F4F000

2103 13.7.4 MCTP control protocol version information

2104 MCTP implementations that follow this particular specification shall return the following version
2105 information in the response to the Get MCTP Version Support message when the Message Type
2106 parameter in the request is set to 0x00 (return MCTP control protocol version information).

2107 The Version Number Entry 1 field shall be used to indicate backward compatibility with Version 1.0 of the
2108 base specification Control Protocol as:

2109 1.0 [Major version 1, minor version 0, any update version, no alpha]

2110 This is reported using the encoding as: 0xF1F0FF00

2111 The Version Number Entry 2 field shall be used to indicate backward compatibility with Version 1.1 of the
2112 base specification Control Protocol as:

2113 1.1 [Major version 1, minor version 1, any update version, no alpha]

2114 This is reported using the encoding as: 0xF1F1FF00

2115 The Version Number Entry 3 field shall be used to indicate backward compatibility with Version 1.2 of the
2116 base specification Control Protocol as:

2117 1.2 [Major version 1, minor version 2, any update version, no alpha]

2118 This is reported using the encoding as: 0xF1F2FF00

2119 The Version Number Entry 4 field shall be used to indicate backward compatibility with Version 1.2 of the
2120 base specification Control Protocol as:

2121 1.3 [Major version 1, minor version 3, any update version, no alpha]

2122 This is reported using the encoding as: 0xF1F3FF00

2123 The version of the MCTP base specification Control Protocol for this specification shall be reported in
2124 Version Number Entry 5 as:

2125 **1.4.0** [Major version 1, minor version 4, update version 0, no alpha]

2126 This is reported using the encoding as: 0xF1F4F000

2127

2128 13.8 Get Message Type Support

2129 The Get Message Type Support command enables management controllers to discover the MCTP
2130 control protocol capabilities supported by other MCTP endpoints, and get a list of the MCTP message
2131 types that are supported by the endpoint. The request and response parameters for this message are
2132 listed in Table 19.

2133 The response to this command may be specific according to which bus the request was received over
2134 (that is, a device that supports a given message type may not support that message type equally across
2135 all buses that connect to the device).

2136

Table 19 – Get Message Type Support message

	Byte	Description
Request data	–	–
Response data	1	Completion Code.
	2	MCTP Message Type Count. One-based. Number of message types in addition to the MCTP control message type that is supported by this endpoint
	(3:N)	List of Message Type numbers. One byte per number. See Table 3 and MCTP Message Types table in MCTP ID .

2137 13.9 Get Vendor Defined Message Support

2138 13.9.1 Overview

2139 The Get Vendor Defined Message Support operation enables management controllers to discover
2140 whether the endpoint supports vendor-defined messages, and, if so, the vendors or organizations that
2141 defined those messages. The format and definition of the request and response parameters for this
2142 message is given in Table 20.

2143

2144

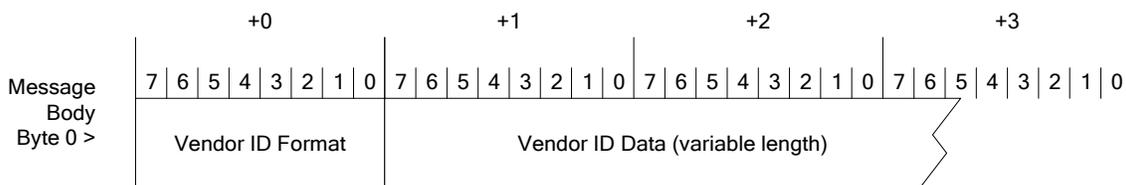
Table 20 – Get Vendor Defined Message Support message

	Byte	Description
Request data	1	Vendor ID Set Selector Indicates the specific capability set requested. Indices start at 0x00 and increase monotonically by 1. If the responding endpoint has one or more capability sets with indices greater than the requested index, it increments the requested index by 1 and returns the resulting value in the response message. The requesting endpoint uses the returned value to request the next capability set.
Response data	1	Completion Code
	2	Vendor ID Set Selector 0xFF = no more capability sets.
	Var	Vendor ID A structured field of variable length that identifies the vendor ID format (presently PCI or IANA) and the ID of the vendor that defined the capability set. The structure of this field is specified in Figure 22 – Structure of Vendor ID field for Get Vendor Defined capabilities message.
	2 bytes	16-bit numeric value or bit field, as specified by the vendor or organization identified by the vendor ID. This value is typically used to identify a particular command set type or major version under the given vendor ID.

2145 **13.9.2 Vendor ID formats**

2146 Figure 22 shows the general structure of Vendor ID fields used in this specification. The first byte of the
 2147 field contains the Vendor ID Format, a numeric value that indicates the definition space and format of the
 2148 ID. The remainder of the field holds the Vendor ID Data with content and format as specified in Table 21.

2149 The MCTP management controller or management device can pick which format is best suited for the
 2150 device. In general, if the device does not already have an existing vendor ID that matches one of the
 2151 specified formats, it is recommended that the IANA enterprise number format be used.



2152

2153 **Figure 22 – Structure of Vendor ID field for Get Vendor Defined capabilities message**

2154

Table 21 – Vendor ID formats

Vendor ID Format Name	Vendor ID Format	Vendor ID Data Length	Description
PCI Vendor ID	0x00	2	16-bit Unsigned Integer. The PCI 2.3 specifications state the following about the PCI vendor ID: “This field identifies the manufacturer of the device. Valid vendor identifiers are allocated by the PCI SIG to ensure uniqueness. 0xFFFF is an invalid value for the Vendor ID.” However, for MCTP this value may be used for identifying aspects other than the manufacturer of the device, such as its use in the Vendor Defined – PCI message type, where it identifies the vendor or organization that defined a particular set of vendor-defined messages. Thus, in some uses, the ID may or may not correspond to the PCI ID for the manufacturer of the device.
IANA Enterprise Number	0x01	4	32-bit Unsigned Integer. The IANA enterprise number for the organization or vendor expressed as a 32-bit unsigned binary number. For example, the enterprise ID for the DMTF is 412 (decimal) or 0x0000_019C expressed as a 32-bit hexadecimal number. The enterprise number is assigned and maintained by the Internet Assigned Numbers Authority, www.iana.org, as a means of identifying a particular vendor, company, or organization.

2155 **13.10 Resolve Endpoint ID**

2156 This command is sent to the bus owner to resolve an EID into the physical address that shall be used to
 2157 deliver MCTP messages to the target endpoint. The command takes an EID as an input parameter in the
 2158 request and returns the EID and the physical address for routing to that EID (if any) in the response. The
 2159 response data will also indicate if no mapping was available.

2160 An endpoint knows the physical address of the bus owner by keeping track of which physical address
 2161 was used when the endpoint received its EID assignment through the Set Endpoint ID command. The
 2162 endpoint can send this command to the bus owner using the null destination EID value. This eliminates
 2163 the need for the endpoint to also keep track of the EID of the bus owner. The request and response
 2164 parameters are specified in Table 22.

2165 **Table 22 – Resolve Endpoint ID message**

	Byte	Description
Request data	1	Target Endpoint ID This is the EID that the bus owner is being asked to resolve.
Response data	1	Completion Code
	2	Bridge Endpoint ID This is the EID for the endpoint that is providing the bridging server (if any) that is required to access the target endpoint. If the EID being returned matches the same value as the target EID, it indicates that there is no bridging function that is required to access the target endpoint (that is, the target EID is local to the bus that the Resolve Endpoint ID request was issued over).
	3:N	Physical Address.

	Byte	Description
		The size of this field is dependent on the particular MCTP physical transport binding used for the bus that this data is being provided for. The size and format of this field is defined as part of the corresponding physical transport binding specification.

2166 **13.11 Allocate Endpoint IDs**

2167 Bus owners are responsible for allocating pools of EIDs to MCTP bridges that are lower in the bus
2168 hierarchy. This is done using the Allocate Endpoint IDs command. The EID for the bridge itself is
2169 assigned separately and is *not* part of the pool given with this command.

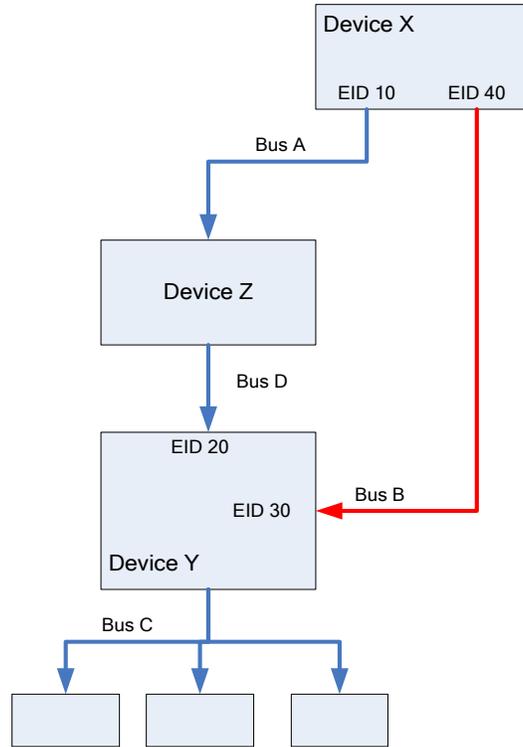
2170 The bus owner will typically use this command as part of the EID assignment process for a bus. When a
2171 device has been assigned an EID using the Set Endpoint ID command, the response to that command
2172 indicates whether the endpoint supports an EID pool. If the device indicates that it supports an EID pool,
2173 the bus owner can then issue the Allocate Endpoint IDs command to supply the pool of EIDs to the
2174 device.

2175 NOTE: The Allocate Endpoint IDs command can also cause a bridge to rebuild its routing table. See 13.12.3 for
2176 more information.

2177 When an EID or EID pool that was previously allocated becomes unused (for example, due to a hot-swap
2178 removal), the bus owner shall reclaim the endpoint's EID or EID pool allocation. See 8.18 for additional
2179 details.

2180 Referring to Figure 23, there is a potential race condition with handling EID allocation. In the scenario
2181 shown in this figure, it is possible that device X and device Z might both be assigning EIDs to device Y at
2182 the same time. This also means that, unless steps are taken, device Z could allocate endpoints to device
2183 Y only to have this overwritten by a set of endpoints assigned by device X.

2184 To prevent this, the Allocate Endpoint IDs command is only accepted from the "first" bus that provides the
2185 EID pool to the device. If another bus owner attempts to deliver an EID pool through another bus, the
2186 request will be rejected unless an intentional over-ride is done.



2187

2188

Figure 23 – EID Pools from multiple bus owners

2189 The Allocate Endpoint IDs message fields are described in Table 23.

2190

Table 23 – Allocate Endpoint IDs message

	Byte	Description
Request data	1	Operation Flags: [7:2] – reserved. [1:0] – Operation: 00b = Allocate EIDs. Submit an EID pool allocation. Do not force allocation. This enables the allocation to be rejected if the bridge has already received its EID pool from another bus. (See additional information in the following clauses.) 01b = Force allocation. Force bridge to accept this EID pool regardless of whether it has already received its EID pool from another bus. This shall also cause a bridge to rebuild its routing tables. See 13.12.3 for more information. 10b = Get allocation information Return the response parameters without changing the present allocation. This can be used to query information on the dynamic pool of EIDs presently allocated to the Endpoint, if any. If this operation is selected, the Number of Endpoint IDs and Starting Endpoint ID parameters in the request shall be ignored. 11b = Reserved
	2	Number of Endpoint IDs (Allocated Pool Size) Specifies the number of EIDs in the pool being made available to this Endpoint

	Byte	Description
		Specifying a count of 0x00 shall be legal. If 0x00 is accepted or forced (and the bridge lacks a static EID pool) no EIDs shall be available for distribution by the particular bridge.
	3	Starting Endpoint ID Specifies the starting EID for the range of EIDs being allocated in the pool. When multiple EIDs are provided, the IDs are sequential starting with this value as the first EID in the range.
Response data	1	Completion Code An error completion code (ERROR_INVALID_DATA should be returned) shall be returned if the number of EIDs being allocated (Number of Endpoint IDs) exceeds the Dynamic Endpoint ID Pool size. (This error condition does not apply to when the number of endpoint IDs passed in the request is 0x00).
	2	[7:2] – reserved [1:0] – 00b = Allocation was accepted. In the case that the bridge has a completely static EID pool, the bridge should not track which bus has sourced the command and shall accept the allocation if the Number of Endpoint IDs (Allocated Pool Size) is 0x00. 01b = Allocation was rejected. The Allocate Endpoint IDs command is accepted only from the “first” bus that provides the EID pool to the device. If another bus owner attempts to deliver an EID pool through another bus, the request will be rejected unless an intentional over-ride is done. (The rationale for this behavior is explained in the text of this clause.) 10b, 11b = reserved
	3	Endpoint ID Pool Size (Dynamic) This value is the size of the EID pool used by this endpoint. This is the size of the dynamic EID pool that the bridge can use to assign EIDs or EID pools to other endpoints or bridges. It does not include the count of any additional static EIDs that the bridge may maintain. See 8.18.3 for more information.
	4	First Endpoint ID This field specifies the first EID assigned to the pool for this endpoint. The value is 0x00 if there are no EIDs assigned to the pool.

2191 **13.12 Routing Information Update**

2192 **13.12.1 General**

2193 The Routing Information Update message is used by a bus owner to give routing information to a bridge
2194 for the bus on which the message is being received.

2195 Because the physical address format is based on the bus over which the request is delivered, the bus
2196 owner shall use the medium-specific physical address format for the addresses sent using this command.

2197 An MCTP bridge may be sent more than one instance of this command to transfer the update information.
2198 An integral number of routing information update entries shall be provided in the command (that is,
2199 routing information update entries cannot be split across instances of the command).

2200 **13.12.2 Adding and replacing entries**

2201 The recipient of this command shall check to see whether the information in the request corresponds to
 2202 the EID for an existing entry for the bus over which the command was received. If so, it shall replace that
 2203 entry with the new information. If an entry for a given EID or EID range does not already exist, it shall
 2204 create new entries for the given EIDs. In some cases this may require the bridge to split existing entries
 2205 into multiple entries.

2206 NOTE: A bus owner is only allowed to update entries that correspond to its bus. For each routing table entry that
 2207 was created or updated through the Routing Information Update message, the bridge shall keep track of which bus it
 2208 received the Routing Information Update from. This is necessary so that when a Routing Information Update is
 2209 received from a particular bus, the bridge only updates entries that correspond to entries that were originally given to
 2210 it from that bus.

2211 **13.12.3 Rebuilding routing tables**

2212 A bridge that receives and accepts the Allocate Endpoint IDs command with the “Force Allocation” bit set
 2213 (1b) shall clear out and rebuild its routing table information. The bridge shall issue commands to reassign
 2214 EIDs and re-allocate EID pools to all downstream devices. The request and response parameters are
 2215 specified in Table 24, and format information is provided in Table 25.

2216 **Table 24 – Routing Information Update message**

	Byte	Description
Request data	1	Count of update entries (1-based)
	see text	One or more update entries, based on the given count, as illustrated in Table 25
Response data	1	Completion Code 0x80 = Insufficient space to add requested entries to internal routing table

2217 **Table 25 – Routing Information Update entry format**

Byte	Description
1	[7:4] – reserved [3:0] – Entry Type: 00b = entry corresponds to a single endpoint that is not serving as an MCTP bridge 01b = entry reflects an EID range for a bridge where the starting EID is the EID of the bridge itself and additional EIDs in the range are routed by the bridge 10b = entry is for a single endpoint that is serving as an MCTP bridge 11b = entry is an EID range for a bridge, but does not include the EID of the bridge itself
2	[7:0] Size of EID Range. The count of EIDs in the range.
3	First EID in EID Range. The EID Range is sequential (for example, if the size of the EID Range is 3 and the First EID value given in this parameter is 21, the Entry covers EIDs 21, 22, and 23).
4:N	Physical Address. The size and format of this field is defined as part of the corresponding physical transport binding specification for the bus that this data is being provided for.

2218 **13.13 Get Routing Table Entries**

2219 This command can be used to request an MCTP bridge or bus owner to return data corresponding to its
 2220 present routing table entries. This data is used to enable troubleshooting the configuration of routing
 2221 tables and to enable software to draw a logical picture of the MCTP network. More than one instance of
 2222 this command will typically need to be issued to transfer the entire routing table content.

2223 An integral number of routing table entries shall be provided in the response to this command (that is,
 2224 routing table entries cannot be split across instances of the command). The request and response
 2225 parameters are specified in Table 26, and format information is provided in Table 27.

2226 **Table 26 – Get Routing Table Entries message**

	Byte	Description
Request data	1	Entry Handle (0x00 to access first entries in table)
Response data	1	Completion Code
	2	Next Entry Handle (Use this value to request the next set of entries, if any.) If the routing table data exceeds what can be carried in a single MCTP control response. 0xFF = No more entries
	3	Number of routing table entries being returned in this response
	4:N	One or more routing table entries, formatted per Table 27. This field will be absent if the number of routing table entries is 0x00.

2227 **Table 27 – Routing Table Entry format**

Byte	Description
1	Size of EID range associated with this entry
2	Starting EID
3	<p>Entry Type/Port Number</p> <p>[7:6] – Entry Type:</p> <ul style="list-style-type: none"> 00b = entry corresponds to a single endpoint that does not operate as an MCTP bridge 01b = entry reflects an EID range for a bridge where the starting EID is the EID of the bridge itself and additional EIDs in the range are routed by the bridge 10b = entry is for a single endpoint that serves as an MCTP bridge 11b = entry is an EID range for a bridge, but does not include the EID of the bridge itself <p>[5] – Dynamic/Static Entry.</p> <p>Indicates whether the entry was dynamically created or statically configured. Note that statically configured routing information shall not be merged with dynamic information when reporting entry information using this command. While an implementation may internally organize its data that way, dynamic and statically configured routing shall be reported as separate entries. Dynamically created entries include entries that were generated from the Routing Information Update command as well as entries that were created as a result of the bridge doing EID assignment and EID pool allocation as a bus owner.</p> <ul style="list-style-type: none"> 0b = Entry was dynamically created 1b = Entry was statically configured <p>[4:0] – Port number</p> <p>This value is chosen by the bridge device vendor and is used to identify a particular bus connection that the physical address for the entry is defined under. In some cases, this number</p>

Byte	Description
	<p>may correspond to an internal “logical” bus that is not directly connected to an external physical bus. Port numbers are required to be static.</p> <p>It is recommended, but not required, that the ports (bus connections) on the bridge be numbered sequentially starting from 0x00. This specification does not define any requirements or recommendations on how port numbers are assigned to corresponding physical connections on a device.</p>
4	Physical Transport Binding Identifier, according to DSP0239.
5	Physical Media Type Identifier, according to DSP0239. This value is used to indicate what format the following physical address data is given in.
6	Physical Address Size. The size in bytes of the following Physical Address field The size is defined as part of the corresponding physical transport binding specification identified by the physical media type identifier.
7:N	Physical Address. The size and format of this field is defined as part of the corresponding physical transport binding specification. The information given in this field is given MSB first. Any unused bits should be set to 0b.

2228 13.14 Prepare for Endpoint Discovery

2229 The Endpoint Discovery message is used to determine if devices on a bus communicate MCTP (see
2230 Table 28). Whether this message is required depends on the particular medium. Currently, this message
2231 may be required only by a particular transport binding, such as PCI Express (PCIe) VDM, because other
2232 bindings such as SMBus/I2C may use other mechanisms for determining this information.

2233 Each endpoint (except the bus owner) on the bus maintains an internal flag called the “Discovered” flag.

2234 The Prepare for Endpoint Discovery command is typically issued as a broadcast Request message on a
2235 given bus that causes each endpoint on the bus to set their respective Discovered flag to the
2236 “undiscovered” state. The flag is subsequently set to the “discovered” state when the Set Endpoint ID
2237 command is received by the endpoint.

2238 An endpoint also sets the flag to the “undiscovered” state at the following times:

- 2239 • Whenever the physical address associated with the endpoint changes or is assigned
- 2240 • Whenever an endpoint first appears on the bus and requires an EID assignment
- 2241 • During operation if an endpoint enters a state that requires its EID to be reassigned
- 2242 • For hot-plug endpoints: After exiting any temporary state where the hot-plug endpoint was
2243 unable to respond to MCTP control requests for more than $T_{RECLAIM}$ seconds (where $T_{RECLAIM}$ is
2244 specified in the physical transport binding specification for the medium used to access the
2245 endpoint). See 8.18.6 for additional information.

2246 Endpoints that have their Discovered flag set to “undiscovered” should use physical addressing for any
2247 received MCTP control message filtering.

2248 Only endpoints that have their Discovered flag set to “undiscovered” will respond to the Endpoint
2249 Discovery message. Endpoints that have the flag set to “discovered” will not respond.

2250 The destination EID for the Prepare for Endpoint Discovery message is set to the Broadcast EID value
2251 (see Table 2) in the request message to indicate that this is a broadcast message. The response
2252 message sets the destination EID to be the ID of the source of the request message, which is typically the
2253 EID of the bus owner. The request and response parameters are specified in Table 28.

2254 The Prepare for Endpoint Discovery message has no effect on existing EID assignments. That is,
 2255 endpoints shall normally retain their EIDs until they are explicitly changed via the Set Endpoint ID
 2256 command, and shall not clear them after getting a “Prepare for Endpoint Discovery” command. (Note that
 2257 endpoints may lose their EIDs under other conditions such as power state changes, etc., as described
 2258 elsewhere in this specification.)

2259 The Endpoint Discovery and Prepare for Endpoint Discovery commands may only be supported on
 2260 particular transport bindings (e.g. MCTP over PCIe Vendor Defined Messaging). If the binding does not
 2261 use this discovery approach (e.g. SMBus/I2C) the endpoint shall return an `ERROR_UNSUPPORTED_CMD`
 2262 completion status for those commands.

2263 **Table 28 – Prepare for Endpoint Discovery message**

	Byte	Description
Request data	–	–
Response data	1	Completion Code

2264 **13.15 Endpoint Discovery**

2265 This command is used to discover endpoints that have their Discovered flag set to “undiscovered”. Only
 2266 endpoints that have their Discovered flag set to “undiscovered” will respond to this message. Endpoints
 2267 that have the flag set to “discovered” will not respond.

2268 This message is typically sent as a Broadcast Request message by the bus owner using the Broadcast
 2269 EID as the destination EID, though for testing purposes endpoints shall also accept and handle this
 2270 command as a non-broadcast Request. Additionally, the request may be sent as a datagram, depending
 2271 on the transport binding requirements. The request and response (if any) parameters are specified in
 2272 Table 29.

2273 **Table 29 – Endpoint Discovery message**

	Byte	Description
Request data	–	–
Response data	1	Completion Code

2274 **13.16 Discovery Notify**

2275 This message is available for use as a common message for enabling an endpoint to announce its
 2276 presence to the bus owner. This will typically be used as part of the endpoint discovery process when an
 2277 MCTP device is hot-plugged onto or becomes powered up on an MCTP bus.

2278 Whether and how this message is used for endpoint discovery depends on the particular physical
 2279 transport binding specification. For example, the SMBus/I2C transport binding does not use this message
 2280 for an endpoint to announce itself because it takes advantage of mechanisms that are already defined for
 2281 SMBus.

2282 This message should only be sent from endpoints to the bus owner for the bus that the endpoint is on so
 2283 it can notify the bus owner that the endpoint has come online and may require an EID assignment or
 2284 update. Additionally, the request may be sent as a datagram, depending on the transport binding
 2285 requirements. The request and response (if any) parameters are specified in Table 30.

2286

Table 30 – Discovery Notify message

	Byte	Description
Request data	–	–
Response data	1	Completion Code

2287 **13.17 Get Network ID**

2288 The Get Network ID command returns a universally unique identifier (UUID), also referred to as a globally
 2289 unique ID (GUID), for a given MCTP network. Typically this command is sent to the topmost MCTP bus-
 2290 owner since the topmost bus-owner has this knowledge. A Network ID is required for add-in MCTP
 2291 networks (For example, an MCTP Network on an add-in card or module). A Network ID is not required for
 2292 a fixed (not add-in) MCTP network provided there is only one network in the system implementation. A
 2293 Network ID is required for fixed MCTP networks when more than one fixed network exists in the system
 2294 implementation and is simultaneously accessible by a common entity such as system software.

2295 The format of the ID follows the byte (octet) format specified in [RFC4122](#). [RFC4122](#) specifies four
 2296 different versions of UUID formats and generation algorithms suitable for use for a device UUID in IPMI.
 2297 These are version 1 (0001b) “time based”, and three “name-based” versions: version 3 (0011b) “MD5
 2298 hash”, version 4 (0100b) “Pseudo-random”, and version 5 “SHA1 hash”. The version 1 format is
 2299 recommended. However, versions 3, 4, or 5 formats are also allowed. A device UUID should never
 2300 change over the lifetime of the device. The request and response parameters are specified in Table 16.

2301

Table 31 – Get Network ID message format

	Byte	Description
Request data	–	–
Response data	1	Completion Code
	2:17	Network ID bytes 1:16, respectively (see Table 17)

2302 The individual fields within the UUID are stored most-significant byte (MSB) first per the convention
 2303 described in [RFC4122](#). See Table 17 for an example format.

2304 **13.18 Query Hop**

2305 This command can be used to query a bridge to find out whether a given EID shall be accessed by going
 2306 through that bridge, and if so, whether yet another bridge shall be passed through in the path to the
 2307 endpoint, or if the endpoint is on a bus that is directly connected to the bridge.

2308 The command also returns the information about the transmission unit information that the bridge
 2309 supports in routing to the given target endpoint from the bus that the request was received over. See 9.5
 2310 for more information.

2311 **NOTE** The physical transport binding for MCTP may place additional requirements on the physical packet sizes
 2312 that can be used to transfer MCTP packet payloads, such as requiring that physical packet sizes be in 32-byte or 64-
 2313 byte increments, or particular power of 2 increments (for example, 128, 256, 512, and so on).

2314 The request and response parameters are specified in Table 32.

2315

Table 32 – Query Hop message

	Byte	Description
Request data	1	Target Endpoint ID

	Byte	Description
		0x00, 0xFF = reserved. (An ERROR_INVALID_DATA completion code shall be returned.)
	2	Message type for which transmission unit information is being requested. Use the MCTP control message type number unless another message type is of interest.
Response data	1	Completion Code An ERROR_INVALID_DATA completion code shall be returned if the target EID is not covered by any entry in the bridge's routing table.
	2	EID of the next bridge that is used to access the target endpoint, if any Note: This response depends on which bus port the Query Hop request is received over. If this EID is 00h: The EID is covered by the bridge's routing table, but the target EID does not require access by <i>going through</i> this bridge from the port the request was received over. This response will be returned if the target EID is already local to the bus over which the request is being received. This response is also returned when the target EID is an EID for the bridge itself. If this EID is non-zero <i>and</i> is different than the target EID passed in request: The EID being provided is the EID of the "next bridge" in the path to the target EID. If this EID is equal to the target EID passed in request: The target EID is accessed by going through this bridge and no additional bridges shall be gone through to reach the target.
	3	Message Type. This should return 0xFF to indicate that the information is applicable to all message types that are supported by the bridge. Other values should be interpreted as 0xFF as well.
	4:5	Maximum supported incoming transmission unit size in increments of 16 bytes, starting from the baseline transmission unit size (0x0000 = 64 bytes, 0x0001 = 80 bytes, and so on).
	6:7	Maximum supported outgoing transmission unit size in increments of 16 bytes, starting from the baseline transmission unit (0x0000 = 64 bytes, 0x0001 = 80 bytes, and so on). The responder will return whether this transmission unit size is supported for MCTP packets that it transmits for the given message type.

2316 **13.19 Resolve UUID**

2317 This command is used to get information about an endpoint based on its UUID. This command may be
2318 sent from any endpoint to the bus owner. This command takes a UUID as a parameter in the request and
2319 returns a list of EIDs and physical addresses that matches this UUID.

2320 A bus owner that supports this command shall keep in the routing table entries the UUID of each of the
2321 endpoints. The UUID values can be found using a "Get Endpoint UUID" command.
2322 An endpoint knows the physical address of the bus owner by keeping track of which physical address
2323 was used when the endpoint received its EID assignment through the Set Endpoint ID command. The
2324 endpoint can send this command to the bus owner using the null destination EID value. This eliminates
2325 the need for the endpoint to also keep track of the EID of the bus owner. The request and response
2326 parameters are specified in Table 33.

2327

Table 33 – Resolve UUID message

	Byte	Description
Request data	1:16	Requested UUID
	17	Entry Handle (0x00 to access first entries in table)
Response data	1	Completion Code
	2	Next Entry Handle (Use this value to request the next set of entries, if any.) If the EID table data exceeds what can be carried in a single MCTP control response. 0xFF = No more entries
	3	Number of EID entries being returned in this response.
	4:N	One or more routing table entries, formatted per Table 34. This field will be absent if the number of EID entries is 0x00.

2328

Table 34 – Resolve UUID message entry format

Byte	Description
0	EID
1	Physical Transport Binding Type Identifier, according to MCTP ID specification (DSP0239).
2	Physical Media Type Identifier, according to MCTP ID specification (DSP0239). This value is used to indicate what format the following physical address data is given in.
3	Physical Address Size.
4:N	Physical Address.

2329 13.20 Query rate limit

2330 This command can be used to query an EID for its transmit rate limiting capabilities and its receive data
2331 rate requirements.

2332 This command can be used by a message originator to determine the data rate that this EID accepts. The
2333 command can also be used to query the present settings for the EID's transmit data rate capabilities and
2334 present setting.

2335 The request and response parameters are specified in Table 35.

2336

Table 35 – Query rate limit message

	Byte	Description
Request data	-	-
Response data	1	Completion Code
	2:5	Receive information: receive buffer size in bytes.
	6:9	Receive Information: maximum receive data rate limit, in baseline transmission unit packets/sec. A value of 0x0 indicates the receiver is not requesting limiting of the traffic. Note: Unless otherwise specified, it should be assumed that the limit has been defined for communication between two EIDs with the receiver in its most typical modes of operation. The value is not a guarantee. Factors such

	Byte	Description
		as transient loading, and a typical device states may mean the receiver will be temporarily unable to receive at the rate given in this response.
	10:13	Transmit Rate limiter capabilities: Maximum supported rate limit, in baseline transmission unit packets/sec. A value of 0x0 means the device cannot throttle its traffic.
	14:17	Transmit Rate limiter capabilities: Minimum supported rate limit, in baseline transmission unit packets/sec. A value of 0x0 means the device cannot throttle its traffic.
	18:20	Transmit Rate limiter capabilities: Maximum supported burst size.
	21:23	Present Transmit Rate Limit Burst Setting: The maximal burst size allowed to be sent from this EID at one time.
	24:27	Present Setting: EID Maximal Transmit data rate limit, in baseline transmission unit packets/sec. A value of 0x0 means the rate limiter is not active (When Rate Limiting is inactive, the EID will be transmitting at the maximum rate for its present state).
	28	Transmit Rate limiter capabilities: [7:2] – Reserved [1] – Transmit Rate limiting operation capability 0b – Transmit Rate limiting on this EID is applied to requested and non-requested messages together 1b – Transmit Rate limiting on this EID is applied only to non-requested messages [0] - Rate limiting Support on EID 0b – Transmit Rate limiting is not supported 1b – Transmit Rate limiting is supported

2337 **13.21 Request TX rate limit**

2338 This command can be used to configure an EID for its maximal transmit rate limitations settings.

2339 This command shall be used by a data-receiving device to request to configure a transmitting EID for the
 2340 maximal allowed data rate from the transmitting endpoint to that data-receiving EID.

2341 The request and response parameters are specified in Table 36.

2342 **Table 36 – Request TX rate limit message**

	Byte	Description
Request data	1:3	EID transmit maximal burst size in in MCTP packets. This value defines the maximum number of back-to-back consecutive packets that are allowed to be sent from this endpoint, which the receiving EID supports. The term 'back-to-back' means the packets are transmitted with the minimum delay between them. This value shall be set to at least 1 packet to enable rate-limiting. A value of 0 in this field shall be used only to disable rate-limiting.
	4:7	EID Maximal Transmit data rate limit, in baseline transmission unit packets/sec.

	Byte	Description
Response data	1	Completion Code An ERROR_INVALID_DATA shall be returned if the rate limit requested is not supported.
	2:5	EID transmit burst size in MCTP packets. This value defines the presently used maximum total burst size allowed to be sent from this endpoint at one time.
	6:9	EID transmit data rate limit, as presently used, in baseline transmission unit packets/sec.

2343 The response values for EID transmit burst size in MCTP packets, and EID transmit data rate limit, may
 2344 differ from the requested values. This can happen when multiple requests from multiple source EIDs
 2345 received with different request values sharing the same rate limiter. See description in 10.1.6.

2346 The response to this command is sent when the new rate is in effect when a change is performed or
 2347 immediately when no change is done. Following sending a response to Request TX rate limit command
 2348 for the first time from any EID, it is recommended that the endpoint receiving this command will send Get
 2349 Endpoint UUID command to the EID which sent the Request TX rate limit command. This allows any
 2350 device to identify when an endpoint is enumerated with a different EID, in order to properly calculate its
 2351 rate-limiting settings.

2352 13.22 Update rate limit

2353 This command is sent from a transmitter EID to a receiver EID, to update a receiver on any change in the
 2354 transmitter's rate settings, which did not originate from a request from the receiver. This command is sent
 2355 to any connected receive EID which is not the EID which originated the rate change.

2356 The command shall be used only after a change of the EID transmit burst size and/or EID transmit data
 2357 rate limit.

2358 The request and response parameters are specified in Table 38.

2359 **Table 37 – Update rate limit message**

	Byte	Description
Request data	1:3	EID transmit burst size in MCTP packets. This value defines the presently used maximum total burst size allowed to be sent from this endpoint at one time.
	4:7	EID transmit data rate limit, as presently used, in baseline transmission unit packets/sec.
Response data	1	Completion Code

2360 If an error occurred on the transmitter which caused the rate limiting to be set to an unsupported rate, the
 2361 receiver EID shall issue a new Request TX rate limit command to the transmitter EID.

2362 13.23 Query Supported Interfaces

2363 This command can be used to query an endpoint for its MCTP interfaces capabilities.

2364 This command can be used by an MCTP device A to query the different interfaces which are available on
 2365 MCTP device B for communicating MCTP messages between device A and B.

2366 The request and response parameters are specified in Table 38.

2367

Table 38 – Query Supported Interfaces

	Byte	Description
Request data	-	-
Response data	1	Completion Code
	2	Supported Interfaces Count (shall be ≥ 1)
	3	First interface Type (see MCTP physical medium identifiers table in DSP0239)
	4	First interface EID
	...	
	...	
	N-1	Last interface Type (see MCTP physical medium identifiers table in DSP0239)
	N	Last interface EID

2368 **13.24 Query Endpoint Advanced Capabilities**

2369 This command can be used to query an MCTP 1.4.0 (or later version) compatible endpoint, for its
 2370 advanced capabilities support as an endpoint.

2371 This command can be used by an MCTP endpoint A to query the maximal number of concurrent MCTP
 2372 Message transfers which the MCTP device can receive.

2373 The request and response parameters are specified in Table 39.

2374

Table 39 – Query Endpoint Advanced Capabilities

	Byte	Description
Request data	-	-
Response data	1	Completion Code
	2	Maximal number of concurrent MCTP Message Count allowed (shall be ≥ 1)
	3-4	Maximal bytes count in transmission unit size (shall be \geq BTU) Packet errors detection code, if included, is part of the Physical Medium-Specific Trailer, and is not affected by this parameter.

2375 **13.25 Request EID pool**

2376 This command can be used by an MCTP bridge to request allocation of additional EIDs. This command is
 2377 only used by MCTP bridges and should be received by any MCTP bus owner. This command shall only
 2378 be sent to the bus owner which assigned the current pool of EIDs.

2379 The request and response parameters are specified in Table 40.

2380

Table 40 – Request EID pool

	Byte	Description
Request data	1	Number of additional EID requested to be allocated
Response data	1	Completion Code

	Byte	Description
	2	E – represents the number of allocated EIDs. When this parameter is 0, no additional EID can be allocated to the requesting bridge.
	3	N – represents the number of EIDs sections as defined in Table 41. Shall be set to 0 when the number of additional EIDs is 0.
	4- (4+4*N)	These entries are only included when N>0 N entries of EIDs sections as defined in Table 41.

2381

Table 41 – EID section structure

	Byte	Description
Record bytes data	1	NS - Number of EIDs in this section. NS is always > 0.
	2-3	Reserved
	4	Starting EID – represents the first EID number on this section.

2382 **13.26 Release EID pool**

2383 This command can be used by an MCTP bus owner to request an MCTP bridge to return a number of
 2384 allocated EIDs. This command is only used by MCTP bus owners and should be received by any MCTP
 2385 bridge. This command shall only be sent by a bus owner through the same EID which it used to assign
 2386 the pool of EIDs to that bridge. Only un-assigned EIDs can be released by this command. When the
 2387 number of unassigned EIDs is lower than the number of EIDs that were requested in this command, only
 2388 the available EIDs can be returned.

2389 The request and response parameters are specified in Table 39.

2390

Table 42 – Release EID pool

	Byte	Description
Request data	1	Number of requested EIDs to be released
Response data	1	Completion Code
	2	E – represents the number of released EIDs. When this parameter is 0, no additional EID can be released to the requesting MCTP bus owner.
	3	N – represents the number of EIDs sections as defined in Table 41. Shall be set to 0 when the number of additional EIDs is 0.
	4- (4+4*N)	These entries are only included when N>0 N entries of EIDs sections as defined in Table 41.

2391

2392 **13.27 Query Supported Recovery Actions**

2393 This command is used by MCTP bus owners to query the supported recovery methods of an Endpoint
 2394 and/or a downstream bus as described in section 11.

2395

2396 The request and response parameters are specified in Table 43.

2397

Table 43 – Query Supported Recovery Actions

	Byte	Description
Request data	1	Reset Level as described in Table 44
	2	Target type: 0 – MCTP transport protocol 1 – MCTP EID 2 – Port Number as defined in Table 27
	3	Target EID – when Target type is 0 or 1. 0 – Otherwise
	4	0 - When Target type is 0 or 1 [7:5] – 3'b0 [4:0] 5'b0 – When Target type is 0 or 1 Port Number – When Target type is 2
	5	MCTP Message Type – When Target Type is 0 0 - Otherwise
	6:N	PCI Vendor ID – when Message Type = 0x7E IANA Vendor ID – when Message Type = 0x7F 0 – Otherwise
Response data	1	Completion Code
	2	0 – Queried mode is not supported 1 – Queried mode is supported

2398

2399

Table 44 Reset levels

Level	Description
0	Application-level reset as defined in 11.2.1
1	Transport-level reset as defined in 11.2.2
2	Physical Bus reset as defined in 11.2.3
3	Device manageability sub-system reset as defined in 11.2.4
Other	Reserved

2400

2401 **13.28 Recovery Request**

2402 This command is used by MCTP bus owners to issue a recovery method for of an Endpoint and/or a
2403 downstream bus as described in section 11.

2404 The request and response parameters are specified in Table 45.

2405

Table 45 –Recovery Request

	Byte	Description
Request data	1	Reset Level as described in Table 44
	2	Target type: 0 – MCTP transport protocol 1 – MCTP EID 2 – Port Number as defined in Table 27
	3	Target EID – when Target type is 0 or 1. 0 – Otherwise
	4	0 - When Target type is 0 or 1 [7:5] – 3'b0 [4:0] 5'b0 – When Target type is 0 or 1 Port Number – When Target type is 2
	5	MCTP Message Type – When Target Type is 0 0 - Otherwise
	6:N	PCI Vendor ID – when Message Type = 0x7E IANA Vendor ID – when Message Type = 0x7F 0 – Otherwise
Response data	1	Completion Code

2406 13.29 Transport Specific

2407 Transport Specific commands are a range of commands that are available for use by transport binding
 2408 specifications in order to perform additional MCTP Control functions that are defined by a particular
 2409 transport binding. Transport specific commands shall only be addressed to endpoints on the same
 2410 medium. A bridge is allowed to block transport specific commands from being bridged to different media.

2411 The request and response parameters are specified in Table 46.

2412

Table 46 – Transport Specific message

	Byte	Description
Request data	1	MCTP Physical Transport Binding Identifier The ID of the Physical Transport specification that defines the transport specific message. This ID is defined in MCTP physical medium identifiers table in DSP0239 companion document to this specification.
	2	MCTP Physical Media Identifier The ID of the physical medium that the message is targeted for. This ID is defined MCTP physical medium identifiers table in DSP0239 companion document to this specification.
	3:N	Transport specific command data. Defined by the transport binding specification identified by the MCTP Physical Transport Binding Identifier given in byte 1. If the Physical Transport Binding Identifier = Vendor Defined: The first four bytes of data shall be the IANA Enterprise ID for the Vendor. MSB first. See 13.9.2 for the information on the IANA Enterprise ID as used in this specification.

	Byte	Description
Response data	1	Completion Code

2413 **14 Vendor Defined – PCI and Vendor Defined – IANA messages**

2414 The Vendor Defined – PCI and Vendor Defined – IANA message types provide a mechanism for
2415 providing an MCTP message namespace for vendor-specific messages over MCTP.

2416 Note that the timings for MCTP Vendor Defined Messages (upper layer) are out of scope for DSP0236

2417 The PCI and IANA designations refer to the mechanism that is used to identify the vendor or organization
2418 this is specifying the message's functionality and any parametric data or other fields provided in the
2419 message body.

2420 Note that this specification only defines the initial bytes in the message body of these messages, and sets
2421 the requirement that these messages shall follow the requirements set by the MCTP base protocol and
2422 any additional requirements necessary to meet the transport of these messages over a particular
2423 medium, such as path transmission unit limitations.

2424 Otherwise, any other field definitions and higher-level message behavior such as retries, error/completion
2425 codes, and so on, is message type-specific and thus is vendor-specific.

2426

2427 **14.1 Vendor Defined – PCI message format**

2428 For these messages, the MCTP message type is set to the value for "Vendor Defined – PCI" as defined in
 2429 Table 3. The request and response parameters are specified in Table 47.

2430 **Table 47 – Vendor Defined – PCI message format**

	Byte	Description
Request data	1:2	PCI/PCIe Vendor ID. Refer to PCIe . MSB first. This value is formatted per the Vendor Data Field for the PCI Express vendor ID format. See 13.9.2". NOTE: Because the vendor ID format is implied by the command, the Vendor ID Format bytes are not part of this field.
	(3:N+2)	Vendor-Defined Message Body. 0 to N bytes.
Response data	1:2	PCI/PCIe Vendor ID. Refer to PCIe . MSB first.
	(3:M+2)	Vendor-Defined Message Body. 0 to M bytes.

2431 **14.2 Vendor Defined – IANA message format**

2432 For these messages, the MCTP message type is set to the value for "Vendor Defined – IANA" as defined
 2433 in Table 3. The request and response parameters are specified in Table 48.

2434 **Table 48 – Vendor Defined – IANA message format**

	Byte	Description
Request data	1:4	IANA Enterprise ID for Vendor. MSB first. This value is formatted per the Vendor Data Field for the IANA enterprise vendor ID format. See 13.9.2. NOTE: Because the vendor ID format is implied by the command, the Vendor ID Format bytes are not part of this field.
	(5:N+4)	Vendor-Defined Message Body. 0 to N bytes.
Response data	1:4	IANA Enterprise ID for the Vendor. MSB first.
	(5:M+4)	Vendor-Defined Message Body. 0 to M bytes.

2435

ANNEX A (informative)

Notation

2436
2437
2438
2439

2440 A.1 Notations

2441 Examples of notations used in this document are as follows:

- 2442 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets
2443 starting from byte two and continuing to and including byte N. The lowest offset is on
2444 the left, and the highest is on the right.
- 2445 • (6) Parentheses around a single number can be used in message field descriptions to
2446 indicate a byte field that may be present or absent.
- 2447 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range
2448 may be present or absent. The lowest offset is on the left, and the highest is on the
2449 right.
- 2450 • [PCle](#) Underlined, blue text is typically used to indicate a reference to a document or
2451 specification called out in 2, "Normative References", or to items hyperlinked within
2452 the document.
- 2453 • rsvd Abbreviation for Reserved. Case insensitive.
- 2454 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets
2455 are given as zero-based values (that is, the least significant bit [LSb] offset = 0).
- 2456 • [7:5] A range of bit offsets. The most-significant is on the left, and the least-significant is on
2457 the right.
- 2458 • 1b The lower case "b" following a number consisting of 0s and 1s is used to indicate the
2459 number is being given in binary format.
- 2460 • 0x12A A leading "0x" is used to indicate a number given in hexadecimal format.

2461

ANNEX B (informative)

Change log

2462
2463
2464
2465

Version	Date	Description
1.0.0	2009-05-21	
1.1.0	2010-02-19	Updated the glossary and the overview section, including additions for MCTP host interfaces and descriptions of MCTP networks. Added support for MCTP network IDs and the Get Network ID command. Addressed Mantis issue: 0000417. Added text to Clause 1 (Scope) referencing DSP0238, DSP0239 per WG ballot comments.
1.2.0	2013-01-10	Added <i>Resolve UUID</i> command. Clarified use of Control Protocol Version and versioning for OEM commands, <i>Prepare for Endpoint Discovery</i> command, and the <i>Allocate Endpoint IDs</i> command. Clarified requirements on MCTP Control message flags and TO bit use. Changed command requirements to allow an Endpoint to optionally accept or generate Routing Information Update commands. Corrected typographic and formatting errors.
1.2.1	2014-10-09	Corrected misuse of reserved EIDs in figures. Changed document organization to place bridging clauses in a new first level clause "MCTP Bridging". Added clarifications and clause on "Endpoint ID Retention". Added more cross references and clarifications to better identify requirements associated with the <i>Get Endpoint UUID</i> command.
1.3.0	2016-11-24	Added Rate Limiting. Fixed formatting errors and typos. Added Query Supported Interfaces command
1.3.1	2019-11-19	Added acknowledgements list Corrected error in sections 13.7.3 and 13.7.4 Corrected bits field error in Table 14, Table 15 and Table 23
1.3.2	2024-01-30	Clarified dropped packets description in 8.7 Corrected typos in Table 12 and Table 13 Clarified setting of Discovered Flag in 13.4 Added reference to the interface IDs in DSP0239 in 13.23 and 13.24 Editorial fixes. ISO compliancy sections numbering edits and terms definitions edits. Aligned terms between Figure 4 and Table 3 Removed redundant terms and definitions. Updated Table 9 to clarify the EID addressing in control response.
1.3.3	2024-02-13	Added definition for physical layer term in 3.1.12 Updated 13.14 to use address filtering by an undiscovered endpoint Editorial fixes.
1.4.0	2026-1-20	Added control command Query Endpoint Advanced Capabilities Updated Figure 18 and updated cross reference to it in 9.5.4 Revised Message Type field description in command Query Hop Revised description in section 13.14 to indicate that "Prepare for Endpoint Discovery command" is typically issued rather than being required. Add support for a single device with multiple EIDs Typo fix (bit field assignment) in Table 32 Added commands Request EID pool and Release EID pool Added Resiliency and recovery section

		Added commands Query Supported Recovery Actions and Recovery Request Clarified that MCTP Vendor Defined timings are out of scope for DSP0236
--	--	--

2466