



1  
2  
3  
4

**Document Identifier: DSP0237**

**Date: 2020-04-06**

**Version: 1.2.0**

5  
6  
7

# **Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding Specification**

8 **Supersedes: 1.1.0**  
9 **Document Class: Normative**  
10 **Document Status: Published**  
11 **Document Language: en-US**

12

13 Copyright Notice

14 Copyright © 2009, 2017, 2020 DMTF. All rights reserved.

15 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems  
16 management and interoperability. Members and non-members may reproduce DMTF specifications and  
17 documents, provided that correct attribution is given. As DMTF specifications may be revised from time to  
18 time, the particular version and release date should always be 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29 withdrawn or modified after publication, and shall be indemnified and held harmless by any party  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 implementations of 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.

39

40

# CONTENTS

41 Foreword ..... 5

42 Introduction..... 6

43 1 Scope ..... 7

44 2 Normative references ..... 7

45 3 Terms and definitions ..... 7

46 4 Symbols and abbreviated terms ..... 8

47 5 Conventions ..... 10

48 5.1 Reserved and unassigned values ..... 10

49 5.2 Byte ordering..... 10

50 6 MCTP over SMBus/I<sup>2</sup>C transport ..... 11

51 6.1 Terminology ..... 11

52 6.2 Transport binding use with I<sup>2</sup>C..... 11

53 6.3 MCTP packet encapsulation ..... 12

54 6.4 Bridges and packet formatting ..... 13

55 6.5 MCTP support discovery..... 14

56 6.6 Support for fixed-address devices ..... 14

57 6.7 Supported media..... 14

58 6.8 Physical address format for MCTP control messages..... 15

59 6.9 Get endpoint ID medium-specific information..... 15

60 6.10 Bus owner address ..... 15

61 6.11 Bus address assignment ..... 15

62 6.12 SMBus/I<sup>2</sup>C considerations for MCTP messages ..... 19

63 6.13 Fairness arbitration ..... 20

64 6.14 NACK window ..... 21

65 6.15 Fairness arbitration requirements for MCTP bridges..... 22

66 6.16 Fairness arbitration requirements for non-bridge endpoints..... 23

67 6.17 Fairness arbitration timing ..... 24

68 6.18 MCTP packet timing requirements ..... 25

69 6.19 MCTP control message timing requirements..... 27

70 6.20 "Stuck 0" condition handling ..... 29

71 6.21 MCTP over SMBus/I<sup>2</sup>C protocol anti-aliasing ..... 29

72 6.22 Well-known and reserved slave addresses ..... 30

73 6.23 Fixed address allocation ..... 31

74 6.24 Recommended address range allocation for computer systems ..... 32

75 ANNEX A (informative) Notation ..... 36

76 ANNEX B (informative) Change log ..... 37

77

## 78 Figures

79 Figure 1 – MCTP over SMBus/I<sup>2</sup>C packet format ..... 12

80 Figure 2 – Address assignment flow ..... 19

81 Figure 3 – Allowed byte range for first NACK'd byte..... 21

82 Figure 4 – Fairness arbitration timing measurement for SMBus and I<sup>2</sup>C ..... 24

83 Figure 5 – Example system configuration ..... 32

84

85 **Tables**

86	Table 1 – Packet header field descriptions .....	12
87	Table 2 – Supported media .....	15
88	Table 3 – Physical address format .....	15
89	Table 4 – Medium-specific information .....	15
90	Table 5 – Fairness arbitration timing values for 100 kHz SMBus/I <sup>2</sup> C .....	24
91	Table 6 – Fairness arbitration timing values for 400 kHz I <sup>2</sup> C .....	25
92	Table 7 – Fairness arbitration timing values for 1MHz I <sup>2</sup> C .....	25
93	Table 8 – Timing specifications for MCTP packets on SMBus/I <sup>2</sup> C .....	26
94	Table 9 – Timing specifications for MCTP control messages on SMBus .....	27
95	Table 10 – Well-known and reserved slave addresses .....	30
96	Table 11 – Slave address allocation for computer systems .....	34
97		

98

## Foreword

99 The *Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding Specification*  
100 (DSP0237) was prepared by the PMCI Working Group.

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

### 103 **Acknowledgments**

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

#### 105 **Editor:**

- 106 • Yuval Itkin – Mellanox Technologies

#### 107 **Contributors:**

- 108 • Alan Berenbaum – SMSC
- 109 • Patrick Caporale – Lenovo
- 110 • Phil Chidester – Dell Inc
- 111 • Edward Klodnicki - IBM
- 112 • Joe Kozlowski – Dell Inc
- 113 • John Leung - Intel Corporation
- 114 • Eliel Louzoun – Intel Corporation
- 115 • Patrick Schoeller – Hewlett Packard Enterprise
- 116 • Hemal Shah - Broadcom Limited
- 117 • Tom Slaight – Intel Corporation
- 118 • Yi Zeng – Intel Corporation

119

120

## Introduction

121 The Management Component Transport Protocol (MCTP) over SMBus/I2C transport binding defines a  
122 transport binding for facilitating communication between platform management subsystem components  
123 (e.g., management controllers, managed devices) over SMBus/I2C.

124 The *MCTP Base Specification* ([MCTP](#)) describes the protocol and commands used for communication  
125 within, and the initialization of, an MCTP network. The MCTP over SMBus/I2C transport binding definition  
126 in this specification includes a packet format, physical address format, message routing, and discovery  
127 mechanisms for MCTP over SMBus/I2C communications.

128

129  
130

# Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding Specification

## 131 1 Scope

132 This document provides the specifications for the Management Component Transport Protocol (MCTP)  
133 transport binding for SMBus/I<sup>2</sup>C.

## 134 2 Normative references

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

139 DMTF DSP0136, *Alert Standard Format Specification 2.0*,  
140 <https://www.dmtf.org/sites/default/files/standards/documents/DSP0136.pdf>

141 DMTF, DSP0236, *Management Component Transport Protocol (MCTP) Base Specification 1.3*, MCTP,  
142 [http://www.dmtf.org/sites/default/files/standards/documents/DSP0236\\_1.3.pdf](http://www.dmtf.org/sites/default/files/standards/documents/DSP0236_1.3.pdf)

143 DMTF, DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.4*, MCTP\_ID,  
144 [http://www.dmtf.org/sites/default/files/standards/documents/DSP0239\\_1.4.pdf](http://www.dmtf.org/sites/default/files/standards/documents/DSP0239_1.4.pdf)

145 IPMI Consortium, *Intelligent Platform Management Interface Specification, Second Generation 2.0*, 2006,  
146 [ftp://download.intel.com/design/servers/ipmi/IPMIv2\\_0rev1\\_0.pdf](ftp://download.intel.com/design/servers/ipmi/IPMIv2_0rev1_0.pdf)

147 ISO/IEC Directives, Part 2, *Rules for the structure and drafting of International Standards*,  
148 <http://isotc.iso.org/livelink/livelink?func=ll&objId=4230456&objAction=browse&sort=subtype>

149 NXP Semiconductors, *I<sup>2</sup>C-bus specification and user manual*, 4 April 2014  
150 [http://www.nxp.com/documents/user\\_manual/UM10204.pdf](http://www.nxp.com/documents/user_manual/UM10204.pdf)

151 System Management Bus (SMBus) Specification, version 3.1 19-Mar-2018  
152 [http://smbus.org/specs/SMBus\\_3\\_1\\_20180319.pdf](http://smbus.org/specs/SMBus_3_1_20180319.pdf)

## 153 3 Terms and definitions

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

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

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

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

## 167 **4 Symbols and abbreviated terms**

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

### 169 **4.1**

#### 170 **ACK**

171 acknowledge

### 172 **4.2**

#### 173 **ARP**

174 Address Resolution Protocol

### 175 **4.3**

#### 176 **ASF**

177 Alert Standard Format

### 178 **4.4**

#### 179 **BMC**

180 baseboard management controller

### 181 **4.5**

#### 182 **EEPROM**

183 Electrically Erasable Programmable Read-Only Memory

### 184 **4.6**

#### 185 **EID**

186 endpoint identifier

### 187 **4.7**

#### 188 **I<sup>2</sup>C**

189 Inter-Integrated Circuit

### 190 **4.8**

#### 191 **I/O**

192 input/output

### 193 **4.9**

#### 194 **IPMB**

195 Intelligent Platform Management Bus

### 196 **4.10**

#### 197 **IPMI**

198 Intelligent Platform Management Interface

### 199 **4.11**

#### 200 **kHz**

201 kilohertz



202	<b>4.12</b>
203	<b>LSb</b>
204	least significant bit
205	<b>4.13</b>
206	<b>LSB</b>
207	least significant byte
208	<b>4.14</b>
209	<b>max</b>
210	maximum
211	<b>4.15</b>
212	<b>MCTP</b>
213	Management Component Transport Protocol
214	<b>4.16</b>
215	<b>min</b>
216	minimum
217	<b>4.17</b>
218	<b>ms</b>
219	millisecond
220	<b>4.18</b>
221	<b>MSB</b>
222	most significant byte
223	<b>4.19</b>
224	<b>MTU</b>
225	Maximum Transmission Unit
226	<b>4.20</b>
227	<b>NACK</b>
228	not acknowledge
229	<b>4.21</b>
230	<b>PCI</b>
231	peripheral component interconnect
232	<b>4.22</b>
233	<b>PCIe®</b>
234	PCI Express™
235	<b>4.23</b>
236	<b>PEC</b>
237	packet error code
238	<b>4.24</b>
239	<b>PMCI</b>
240	Platform Management Component Intercommunications

241 **4.25**  
242 **PSA**  
243 persistent slave address

244 **4.26**  
245 **rsvd**  
246 reserved (not case sensitive)

247 **4.27**  
248 **SCL**  
249 serial clock

250 **4.28**  
251 **SDA**  
252 serial data

253 **4.29**  
254 **sec**  
255 second

256 **4.30**  
257 **SEEPROM**  
258 serial EEPROM

259 **4.31**  
260 **SMBus**  
261 System Management Bus

262 **4.32**  
263 **UDID**  
264 unique device identifier

## 265 **5 Conventions**

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

### 267 **5.1 Reserved and unassigned values**

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

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

### 272 **5.2 Byte ordering**

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

## 275 6 MCTP over SMBus/I<sup>2</sup>C transport

276 The MCTP over SMBus/I<sup>2</sup>C transport binding defines how MCTP packets are delivered over a physical  
277 SMBus or I<sup>2</sup>C medium using SMBus transactions. This includes how physical addresses are used, how  
278 fixed addresses are accommodated, how physical address assignment is accomplished for hot-plug or  
279 other devices that require dynamic physical address assignment, and how MCTP support is discovered.  
280 Timing specifications for bus and MCTP control operations are also given, and a "fairness" protocol is  
281 defined for the purpose of avoiding deadlock and starvation/lockout situations among MCTP endpoints.

282 The binding has been designed to be able to share the same bus as devices communicating using earlier  
283 SMBus/I<sup>2</sup>C management protocols such as Alert Standard Format (ASF) and IPMI, and with vendor-  
284 specific devices using SMBus/I<sup>2</sup>C protocols. The specifications can also allow a given device to  
285 incorporate non-MCTP SMBus functions alongside MCTP. This is described in more detail in 6.21.

### 286 6.1 Terminology

287 According to [SMBus](#), SMBus devices are categorized as follows, where Address Resolution Protocol  
288 (ARP) refers to the SMBus Address Resolution Protocol (a dynamic slave address assignment protocol)  
289 and UDID refers to a "unique device identifier", a 128-bit value that a device uses during the ARP process  
290 to uniquely identify itself. Because these protocols are implemented with command transactions that are  
291 run on top of the SMBus physical specification, it is possible to use these protocols on devices that  
292 support an I<sup>2</sup>C physical interface.

- 293 • **ARP-capable**

294 [SMBus](#) term indicating a device that supports all [SMBus](#) ARP commands with the exception of  
295 the optional Host Notify command. The slave address is assignable. The device supports both  
296 Reset commands.

- 297 • **Fixed and Discoverable**

298 [SMBus](#) term indicating a device supports the Prepare to ARP, directed Get UDID, general Get  
299 UDID, and Assign Address commands. The slave address is fixed; the device will accept the  
300 Assign Address command but will not allow address reassignment. The device supports both  
301 Reset commands.

- 302 • **Fixed - Not Discoverable**

303 [SMBus](#) term indicating a device supports the directed Get UDID command. The slave address  
304 is fixed.

- 305 • **Non-ARP-capable**

306 [SMBus](#) term indicating a device does not support any ARP commands. The slave address is  
307 fixed.

- 308 • **Fixed Address**

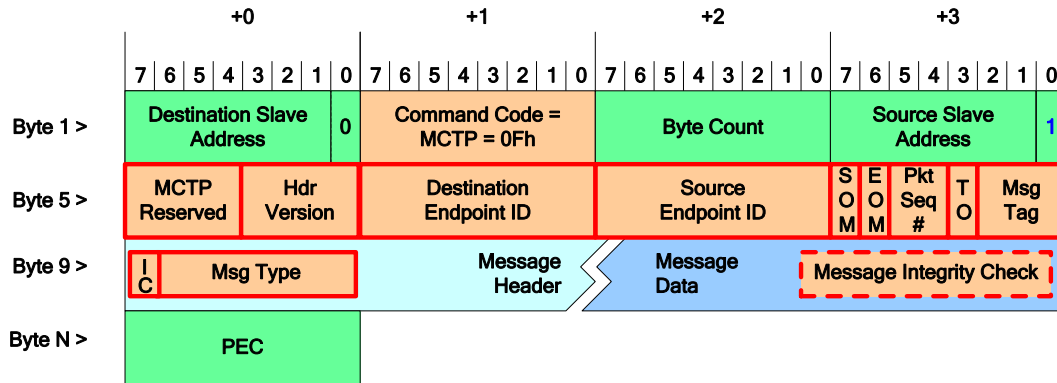
309 For this specification, this term is be used to refer to any device that uses a fixed slave address,  
310 without distinguishing whether it is "Fixed and Discoverable", "Fixed, not Discoverable", or  
311 "Non-ARP-capable".

### 312 6.2 Transport binding use with I<sup>2</sup>C

313 The transport binding defined in this specification has also been designed to be able to work with  
314 standard-mode fast-mode (400 kHz) and Fast-mode Plus (1MHz) I<sup>2</sup>C buses that use 7-bit addressing; 10-  
315 bit addressing is not supported. This binding has not been specified for use with high-speed I<sup>2</sup>C  
316 specifications.

317 **6.3 MCTP packet encapsulation**

318 All MCTP transactions are based on the SMBus Block Write bus protocol. The first 8 bytes make up the  
 319 packet header. The first three fields—Destination Slave Address, Command Code, and Length—map  
 320 directly to SMBus functional fields. The remaining header and payload fields map to SMBus Block Write  
 321 "Data Byte" fields, as indicated in Figure 1. Hence, the inclusion of the Source Slave Address in the  
 322 header is specified by [MCTP](#) rather than [SMBus](#). This is done to facilitate addressing required for  
 323 establishing communications back to the message originator.



324

325 **Figure 1 – MCTP over SMBus/I2C packet format**

326 **Table 1 – Packet header field descriptions**

Byte	Block Write Field(s)	Description
1	Slave Address Wr	[7:1] SMBus Destination Slave Address: The slave address of the target device for the local SMBus link  [0]: SMBus R/W# bit: Shall be set to 0b as all MCTP messages use SMBus write transactions.
2	Command Code	Command Code: SMBus Command Code  All MCTP over SMBus messages use a command code of 0x0F.
3	Byte Count	Byte Count: Byte count for the SMBus Block Write protocol transaction that is carrying the MCTP packet content.  This value is the count of bytes that follow the Byte Count field up to, but not including, the PEC byte. For example, if the MCTP packet payload length (starting with byte 9) is 64 bytes, the value in the Byte Count field would be 69. (The count of 69 accounts for 64 bytes of MCTP packet payload plus the five bytes [bytes 4 through 8, inclusive] that comprise the bytes of the SMBus-specific header and MCTP header that follow the Byte Count field.)

Byte	Block Write Field(s)	Description
4	Data Byte 1	SMBus Source Slave address [7:1] : For the local SMBus link, the slave address of the source device. [0]: This bit shall be set to 1b. The value enables MCTP to be differentiated from IPMI over SMBus and IPMB (IPMI over I <sup>2</sup> C) protocols.
5	Data Byte 2	[7:4] MCTP reserved: This nibble is reserved for definition by the <a href="#">MCTP Base Specification</a> . [3:0] MCTP header version: Set to 0001b for MCTP devices that are conformant to the <a href="#">MCTP Base Specification 1.0</a> and this version of the SMBus transport binding. All other values = Reserved.
6	Data Byte 3	Destination endpoint ID (*)
7	Data Byte 4	Source endpoint ID (*)
8	Data Byte 5	[7] SOM: Start Of Message flag (*) [6] EOM: End Of Message flag (*) [5:4] Packet sequence number (*) [3] Tag Owner (TO) bit (*) [2:0] Message tag (*)
9	Data Byte 6	[7] IC: Integrity Check bit (*) [6:0] Message type (*)
10:N-1	Data Bytes 7:M	Message header and data (*)
N	PEC	Packet error code (PEC): The PEC as defined in the <a href="#">SMBus 2.0 Specification</a> . All MCTP transactions shall include a PEC byte. The PEC byte shall be transmitted by the source and checked by the destination.

(\*) Indicates a field that is defined by the [MCTP Base Specification](#).

327 **6.4 Bridges and packet formatting**

328 As an MCTP packet travels through a bridge from one SMBus/I<sup>2</sup>C port to another, the bridge leaves all  
 329 packet header and message header and data fields alone with the exception of the source and  
 330 destination slave address, which shall be modified to route across the intended bus/link. When an MCTP  
 331 bridge forwards a message from an input port to an output port, it replaces the destination slave address  
 332 with the targeted slave on the destination bus, and replaces the source slave address with the bridge's  
 333 slave address.

334 The MCTP SMBus/I<sup>2</sup>C bridge shall re-calculate the PEC byte to account for changes in the source and  
 335 destination slave address fields.

336 A similar process is used when bridging between different media. The physical addressing and header  
 337 information gets changed by the bridge to match the requirements of the target bus, and any packet-level  
 338 integrity check information is also updated.

## 339 6.5 MCTP support discovery

340 All SMBus devices that support an MCTP endpoint and the SMBus Get UDID command for a particular  
341 SMBus/I<sup>2</sup>C interface (that is, devices with ARP-capable, fixed and discoverable, or fixed-not discoverable  
342 interfaces) are required to have their MCTP support discoverable through the Get UDID command. To do  
343 this, endpoints shall return a value of 1b in bit 5 (the ASF bit) in the Interface field in the Get UDID  
344 command.

345 Once support for ASF has been indicated, an MCTP control message (for example, Get MCTP Version  
346 Support) can be issued to the device to determine whether it supports MCTP. The SMBus command byte  
347 for MCTP packets uses a value that has been allocated by the DMTF for MCTP use and does not overlap  
348 values used for ASF. This enables older devices that indicate ASF support to be queried for MCTP  
349 support without conflict. This is described in more detail in 6.6. Devices that do not support the Get UDID  
350 command will need to have their support for MCTP configured into the bus owner as described in 6.6.

351 I<sup>2</sup>C devices can also support the SMBus protocols and commands for being an ARP-able device that is  
352 also discoverable as an MCTP device. This is required for hot-plug I<sup>2</sup>C devices using MCTP.

## 353 6.6 Support for fixed-address devices

354 MCTP bus owners shall include nonvolatile options to record the addresses used by fixed-address  
355 devices on SMBus/I<sup>2</sup>C buses that they own, and which of those devices support MCTP.

356 For non-MCTP devices, the MCTP bus owner needs this information to know which fixed addresses to  
357 avoid when performing SMBus ARP for the bus. (Alternatively, the bus owner could be configured with a  
358 range of SMBus slave addresses that the bus owner is allowed to allocate from.)

359 For MCTP devices, the bus owner needs this information to perform EID assignment and, if the bus  
360 owner is also an MCTP bridge, routing table initialization and operation.

361 For fixed-address MCTP devices that do not support the Get UDID command (that is, non-ARP-capable  
362 devices), the bus owner needs to also be configured with information that identifies the device as  
363 supporting MCTP.

364 For fixed-address devices that support the [SMBus](#) Get UDID command (that is, devices with ARP-  
365 capable, Fixed and Discoverable, or Fixed-Not Discoverable SMBus interfaces) the bus owner can either  
366 discover whether the device supports MCTP by using the discovery approach described in 6.5, or could  
367 have this information configured at the same time that the slave address information for the fixed-address  
368 device is provided.

369 It is recommended that general-purpose devices that act as MCTP bus owners allow being configured to  
370 support at least 16 different fixed-address devices for each SMBus/I<sup>2</sup>C bus they own. This number would  
371 include both MCTP and non-MCTP devices.

## 372 6.7 Supported media

373 This physical transport binding has been designed to work with the media specified in [DSP0239](#). Table 2  
374 quotes relevant physical media identifiers from [DSP0239](#). In case of any contradiction [DSP0239](#) shall be  
375 used as the normative definition. Use of this binding with other types of physical media is not covered by  
376 this specification. At least one of the physical media identifiers listed in Table 2 shall be supported to  
377 comply with this specification.

378

**Table 2 – Supported media**

Physical Media Identifier	Description
0x01	<a href="#">SMBus</a> 100 kHz compatible
0x02	<a href="#">SMBus</a> + <a href="#">I<sup>2</sup>C</a> 100 kHz compatible
0x03	<a href="#">I<sup>2</sup>C</a> 100 kHz compatible
0x04	<a href="#">I<sup>2</sup>C</a> 400 kHz compatible
0x05	<a href="#">SMBus</a> + <a href="#">I<sup>2</sup>C</a> 1 MHz compatible

379 **6.8 Physical address format for MCTP control messages**

380 The address format shown in Table 3 shall be used for MCTP control commands that require a physical  
 381 address parameter to be returned for a bus that uses this transport binding with one of the supported  
 382 media types listed in 6.7. This includes commands such as the Resolve Endpoint ID, Routing Information  
 383 Update, and Get Routing Table Entries commands.

384

**Table 3 – Physical address format**

Format Size	Layout and Description
1 byte	[7:1] slave address bits [0] 0b

385 **6.9 Get endpoint ID medium-specific information**

386 The medium-specific information as shown in Table 4 shall be used for the medium-specific Information  
 387 field returned in the response to the Get Endpoint ID MCTP control message.

388

**Table 4 – Medium-specific information**

Description
[7:1] reserved
[0] fairness arbitration support (see 6.13) 0b = not supported 1b = supported

389 **6.10 Bus owner address**

390 In order to be the target of the SMBus Notify ARP Master protocol transaction the MCTP bus owner shall  
 391 be configurable to be accessed at the SMBus host slave address. This configuration does not need to be  
 392 used if the bus implementation does not include any MCTP devices that require dynamic address  
 393 assignment of their slave address. For more information, see 6.11.4.

394 The bus owner may use a different, second slave address for all other MCTP communication functions.

395 **6.11 Bus address assignment**

396 This clause describes the configuration, setup, and operation of communication between MCTP  
 397 endpoints using SMBus/I<sup>2</sup>C as the communication medium.

### 398 6.11.1 Slave addresses

399 Each device on SMBus/I<sup>2</sup>C shall have a slave address to be the target of transactions by bus masters.  
400 The MCTP transport protocol solely utilizes Master Write transactions to transfer MCTP packets between  
401 MCTP endpoints. For endpoint "A" to send an MCTP packet to endpoint "B", endpoint A shall master the  
402 bus and issue a Block-Write transaction to the slave address of endpoint B. Similarly, for endpoint B to  
403 send an MCTP packet to endpoint A, it shall master the bus and issue a Block-Write transaction to the  
404 slave address of endpoint A. Thus, bi-directional transfer of MCTP packets requires that both sides of the  
405 communication have slave addresses.

406 Device support for slave addresses can be of two general types: fixed or assignable. Devices with  
407 assignable addresses (also referred to as "ARP-capable" or "ARP-able") can use the [SMBus](#) ARP. The  
408 entity that assigns slave addresses to ARP-able devices is referred to as the "ARP master".

409 A bus can include a mix of fixed-address and ARP-able devices. Most fixed-address devices do not  
410 include a discovery mechanism, and neither [SMBus](#) nor [I<sup>2</sup>C](#) require one. Therefore, for a generic bus  
411 implementation that support ARP-able devices (such as SMBus to PCI/PCIe connectors) the ARP master  
412 needs to know what ranges of addresses are being used for fixed-address devices so that it doesn't give  
413 an ARP-able device an address that conflicts with a fixed-address device.

414 This transport binding allows for non-MCTP devices (both fixed address and ARP-able) to reside on the  
415 same bus segment used for MCTP devices. The use and assignment of slave addresses shall therefore  
416 be compatible with pre-existing devices. To accomplish this, the following approach is used for managing  
417 devices on a bus that supports MCTP.

### 418 6.11.2 Well-known and reserved slave addresses

419 The [SMBus](#) and [I<sup>2</sup>C](#) specifications define certain slave addresses that should either be avoided by  
420 devices or are reserved (not to be used as a general device slave address) because those addresses are  
421 related to functions that are used by MCTP. These addresses are listed in Table 10.

### 422 6.11.3 Fixed address recommendations for device manufacturers

423 MCTP may be used within a typical computer system application where the motherboard/baseboard may  
424 come from one supplier, the chassis from another supplier, and possibly add-in modules from yet  
425 another.

426 Referring to Table 11, it is thus recommended that devices that use fixed addresses and are targeted for  
427 uses that can include baseboard (B), chassis/system (C), and add-in (A) applications are configurable to  
428 cover for at least three different "B" addresses, at least three different "C" addresses, and at least two  
429 different "A" addresses to help avoid address conflicts in those applications.

### 430 6.11.4 Dynamic address assignment (SMBus ARP) support

431 MCTP buses that support connections to standard PCI/PCIe add-in cards are required by the PCI  
432 specifications to support SMBus ARP (be ARP-capable) to allow the devices to be dynamically assigned  
433 addresses to avoid address conflicts and eliminate the need for manual configuration of addresses.  
434 Figure 2 presents an overview of the address assignment process.

### 435 6.11.5 Devices supporting multiple interfaces

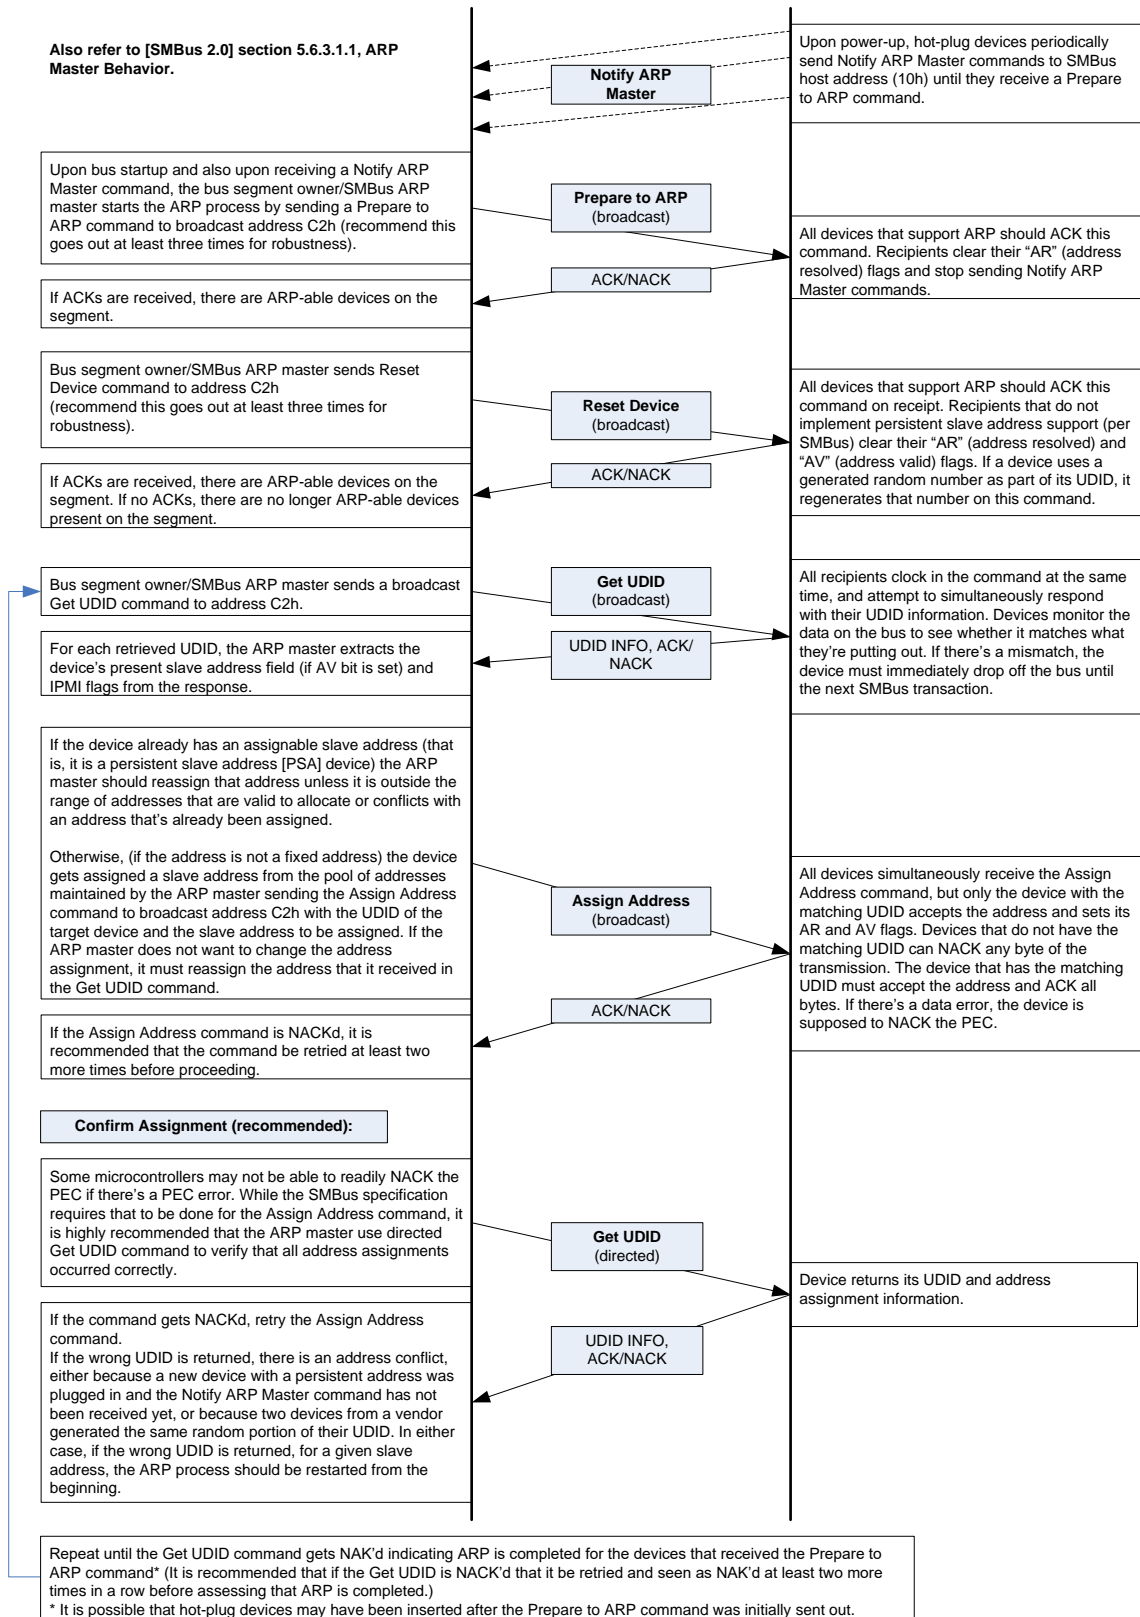
436 Devices that support multiple, separate [SMBus](#) or [I<sup>2</sup>C](#) interfaces where the interfaces are intended to be  
437 connected to the same bus shall meet the following requirements:

- 438 • The interfaces shall be either be ARP-capable or be fixed-address interfaces that are configured  
439 to use a different slave address for each interface.



440           • If the interfaces support SMBus ARP, (as either ARP-able or ARP-enumerable devices) a  
441           different SMBus UDID shall be used for each SMBus ARP-able interface.

442 NOTE   Devices that have internal hardware interfaces that may be implemented as separate blocks but are  
443           designed to share a slave address are not considered to have separate interfaces in this context.



445 **Figure 2 – Address assignment flow**446 **6.11.6 MCTP requirements on SMBus ARP master support**

447 If the bus supports ARP-able devices, MCTP requires that each bus shall have a controller that operates  
448 as the ARP master and assigns slave addresses to all ARP-able devices on the segment. Because the  
449 MCTP bus owner shall know the physical addresses of ARP-able devices that support MCTP, the ARP  
450 master role will typically be handled by the same device that serves as the MCTP bus owner.

451 If a different physical device than the device holding the bus owner functions as the ARP master, there  
452 shall be a mechanism to communicate the address assignment information to the bus owner function.  
453 The mechanism for this is not specified by MCTP.

454 Only one controller is allowed to function as the ARP master for the segment at a given time. The ARP  
455 master function is allowed to *fail-over* or be transferred to another controller. The mechanism for this  
456 capability, if provided, is not specified by MCTP.

457 **6.11.7 Recommendations on ARP master allocation of slave addresses**

458 For PCI and PCI Express™ (PCIe) bus implementations, it is recommended that, by default, the ARP  
459 master only assigns addresses to ARP-able devices from the "B" range. This is because the PCI  
460 slots/connectors themselves are most commonly implemented as part of the board set.

461 Device manufacturers of controllers that function as ARP masters should provide a mechanism to enable  
462 system integrators to either configure which fixed addresses that ARP should avoid, or a pool of non-  
463 conflicting addresses from which ARP can draw.

464 For PCI and PCIe SMBus implementations, the ARP master should be able to assign at least two  
465 addresses for each PCI connector on the segment.

466 **6.11.8 MCTP requirements on hot-pluggable bridges using SMBus**

467 Hot-pluggable MCTP devices that include bridging functionality are required to have static, pre-assigned,  
468 SMBus UDIDs. This is because it is considered a more robust and reliable mechanism than randomly  
469 generated UDIDs, and because it simplifies tracking and managing MCTP device hot-add and hot-  
470 removal.

471 If devices regenerate their UDIDs on hot-plug, the MCTP bus owner/ARP master cannot rely on the UDID  
472 to determine whether a device was newly added to the system. When a hot-plug device includes MCTP  
473 bridging functionality, the bus owner shall be able to allocate the device a range of EIDs from a fixed pool  
474 of IDs. Thus, it is important for the bus owner to be able to determine which devices have been removed  
475 so that any EIDs it had given out can be returned to the pool.

476 It is straightforward for the ARP master to re-enumerate the UDIDs on the bus and determine which  
477 UDIDs (if any) are no longer present (re-enumeration is a natural fallout of the ARP process). If there are  
478 MCTP devices without fixed UDIDs in the mix, however, the bus owner would need to take additional  
479 steps to check to see which devices had already been allocated EIDs to determine by elimination which  
480 ranges, if any, had become freed. With fixed UDID, the bus owner can track which EIDs have been  
481 allocated to which UDIDs and thereby determine which have been freed by a hot swap by just re-  
482 enumerating the UDIDs.

483 **6.12 SMBus/I<sup>2</sup>C considerations for MCTP messages**

484 The following applies to MCTP messages on SMBus regardless of their message type. Note that MCTP  
485 messages require Block Write byte count sizes that exceed limits specified by [SMBus](#). Additional  
486 restrictions on MCTP packets over what the [SMBus](#) and [I<sup>2</sup>C](#) allow are given in 6.3 and 6.18.

### 487 6.12.1 Slave address ACKs/NACKs

488 Per [SMBus](#) and [I<sup>2</sup>C, the NACK of a slave address indicates the physical absence of the device interface.](#)

- 489 • Devices are therefore required to always ACK their slave addresses. This includes ACK'ing  
490 slave addresses used for ARP if the device is ARP-able or ARP-enumerable.
- 491 • An MCTP device *shall* ACK its slave address(es) when the R/W bit on the slave address is 0.

### 492 6.12.2 Clock stretching for non-addressed devices

493 MCTP devices that are monitoring the bus as slaves and do not have a slave address that matches the  
494 transaction shall not clock stretch past the ACK bit for the slave address byte. This requirement only  
495 applies to MCTP packet transactions. It does not apply to non-MCTP-defined messages or transactions,  
496 such as those used for SMBus ARP.

## 497 6.13 Fairness arbitration

### 498 6.13.1 General

499 The following clauses describe an extension to the SMBus/I<sup>2</sup>C arbitration mechanism for device ports that  
500 are used with MCTP. The extensions define a 'fairness' mechanism that helps ensure that ports that are  
501 arbitrating for access to the bus will eventually get access and will not be locked out of access by other  
502 MCTP ports that are using the bus.

503 NOTE Fairness arbitration only applies for messages using the MCTP base protocol. SMBus messages such as  
504 Host Notify are not required to use fairness arbitration.

505 This mechanism works as follows:

- 506 • An MCTP port that wins bus arbitration (per [SMBus](#) or [I<sup>2</sup>C\) for a given transaction shall wait  
507 until it detects a particular bus idle interval before the device can again attempt to arbitrate for  
508 the bus. This is referred to as the device waiting to detect the "FAIR\\_IDLE" condition.](#)
- 509 • Once the port has succeeded in detecting the FAIR\_IDLE condition, it can attempt to get on the  
510 bus and no longer needs to wait to detect the FAIR\_IDLE condition. The port can continue to  
511 attempt to access the bus without waiting for FAIR\_IDLE until the next time the port wins  
512 arbitration. After winning arbitration, the port shall again wait to detect the FAIR\_IDLE condition  
513 before it can attempt to get on the bus.

514 With this approach, all ports that lose arbitration will eventually get a turn at accessing the bus, because  
515 any ports that win arbitration will need to wait until a bus idle interval is detected, while those that have  
516 lost arbitration will not need to wait.

517 For this to work, endpoints shall be able to do two things:

- 518 1) Be able to recognize the FAIR\_IDLE condition. Ports that are waiting to detect a FAIR\_IDLE  
519 condition shall recognize that no other port has made the bus become busy within a particular  
520 window of time ( $T_{IDLE\_WINDOW}$ ) after the bus becomes free.
- 521 2) Ports that have not won arbitration shall be able to issue a START condition soon enough after  
522 the bus becomes free so that a bus busy condition is seen by ports that are waiting to detect a  
523 FAIR\_IDLE condition. To ensure this condition is met, START shall be issued by the port within  
524 a particular window of time ( $T_{START\_WINDOW}$ ) after the bus becomes free.

525 NOTE There is actually no explicit indication in [SMBus](#) or [I<sup>2</sup>C that arbitration has been won. Instead, what the  
526 master detects is that it was able to access the bus and did not have a collision \(lose arbitration\) with  
527 another master. For this specification, this is referred to as \*winning arbitration\*. Because of the way  
528 arbitration works, an MCTP endpoint that is transmitting as a master onto the bus will know that it has won  
529 arbitration if it is able to transmit from the destination slave address byte through the end of the source slave  
530 address byte \(byte 4\) without receiving a collision or NACK.](#)

531 **6.13.2 Deadlock avoidance with fairness arbitration**

532 A device that wins arbitration but is subsequently NACK'd for its write transaction shall return to waiting  
 533 for the FAIR\_IDLE period before it can attempt the transaction again.

534 **6.13.3 Fairness arbitration support**

535 Bridges and endpoints should support fairness arbitration. An endpoint's support for fairness arbitration  
 536 shall be reported through the medium-specific Information field in the response to the Get Endpoint ID  
 537 MCTP control message.

538 **6.13.4 Bus busy sampling requirements for fairness arbitration**

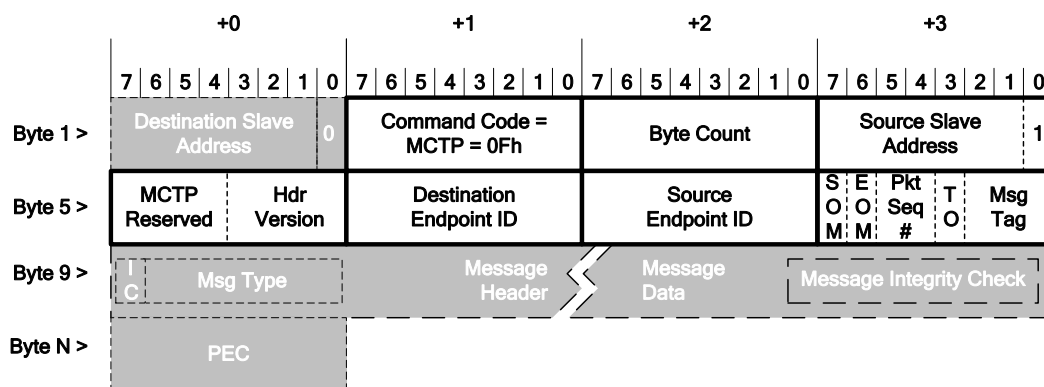
539 It is atypical and unlikely that the bus will go busy and then free again within T\_IDLE\_WINDOW. This is because  
 540 T\_IDLE\_WINDOW is shorter than the time required to send one byte on the bus. Thus, this condition would only  
 541 occur on an error or under a usage of the bus that is not legal within the specifications. Therefore, an  
 542 implementation is not required to continuously check the bus busy status during the entire duration of  
 543 T\_IDLE\_WINDOW (though this is recommended). An implementation is allowed to check the bus busy status  
 544 only at the conclusion of the T\_IDLE\_WINDOW interval that is measured by the device.

545 **6.14 NACK window**

546 An endpoint/bridge is required to NACK an incoming packet if the device does not have input buffer  
 547 space available for the packet. For the NACK to be recognized by the transmitter as the NACK for a  
 548 packet retry, the first NACK bit shall be issued no earlier than byte two (that is, the Command Code byte)  
 549 and no later than byte 8 (the MCTP flags byte). These bytes are represented by the bold outlined bytes in  
 550 Figure 3 below. After the first NACK has been issued any subsequent bytes that are received for the  
 551 packet shall also be NACK'd until a START, STOP, or bus free condition is detected.

552 An endpoint/bridge that NACKs a packet shall continue to NACK any remaining bytes for the transaction  
 553 until it recognizes the next START or STOP condition on the bus.

554



555

556 **Figure 3 – Allowed byte range for first NACK'd byte**

557

## 558 6.15 Fairness arbitration requirements for MCTP bridges

559 MCTP bridges that support fairness arbitration shall meet the following requirements:

- 560 • The bridge shall support FAIR\_IDLE detection and implement the corresponding fairness policy  
561 separately for each port on the bridge.
- 562 • Upon device power up or initialization, a port does not need to detect a FAIR\_IDLE condition  
563 before first attempting to access the bus.
- 564 • A bridge that loses arbitration when attempting to transmit shall continue to retry the transaction  
565 when the bus becomes free for up to PN2 retries (see Table 8). If the retry limit is reached, the  
566 bridge shall drop the packet data.
- 567 • A bridge that receives a NACK when attempting to transmit to a given physical address shall  
568 continue to retry the transaction when the bus becomes free for up to PN2 retries. The bridge  
569 will return to attempting to arbitrate for the bus as described in the preceding requirement,  
570 restarting its number of arbitration retries. If the retry limit is reached, the bridge shall drop the  
571 packet data.
- 572 • An MCTP bridge shall provide dedicated input buffer space per port. The minimum input buffer  
573 size is large enough to store one full baseline MTU-sized MCTP packet. It is recommended, but  
574 not required, that a bridge also implement a dedicated output buffer per port, sized to store at  
575 least one full baseline MTU-sized MCTP packet.
- 576 • If the MCTP bridge is the target of an MCTP packet and it does not have enough buffer space in  
577 its input buffer to store the full packet, it shall NACK the packet. If the bridge has an output  
578 packet to transmit on that same port, it shall be able to issue a START within  $T_{START\_WINDOW}$  after  
579 issuing the retry NACK.
- 580 • A bridge is required to drop a received packet if it finds that the packet error code (PEC) byte for  
581 the transaction is incorrect.
- 582 • An MCTP bridge is not allowed to perform "connected" transactions where the decision to ACK  
583 or NACK an incoming packet is dependent on the bridge's ability to acquire the destination bus  
584 prior to accepting the packet.
- 585 • MCTP bridges are required to implement "store and forward" packet processing. That is, once a  
586 bridge has accepted a packet for routing, it shall retain that packet until it can successfully  
587 transmit it onto the target bus (except when running out of retries when trying to access the  
588 target bus, or upon receiving a packet for a bus that is unavailable or an endpoint that is not  
589 present.)
- 590 • A bridge cannot make the acceptance of a receive packet on its upstream port (port that  
591 connects to a bus that is not owned by the bridge itself) conditional on its ability to transmit a  
592 packet on its upstream port. This requirement does not apply to a downstream port on a bridge  
593 (that is, a downstream port may elect to NACK an incoming packet to allow the bridge to  
594 transmit from that port). This requirement is to help avoid deadlock situations if a bridge is  
595 required to route a packet back onto the bus from which the packet came.
- 596 • A bridge that receives a NACK while it is performing a Master Write operation is not required to  
597 immediately conclude the Master Write operation and drop off the bus. The bridge may continue  
598 the write operation through its conclusion. In either case, the master shall always conclude its  
599 transaction with a STOP condition, unless some other device on the bus first produces a  
600 START or STOP condition. The latter situation is an erroneous condition on the bus, but bridges  
601 shall be able to handle it. Devices shall always recognize START and STOP conditions  
602 regardless of the transaction or bit position on which they occur.

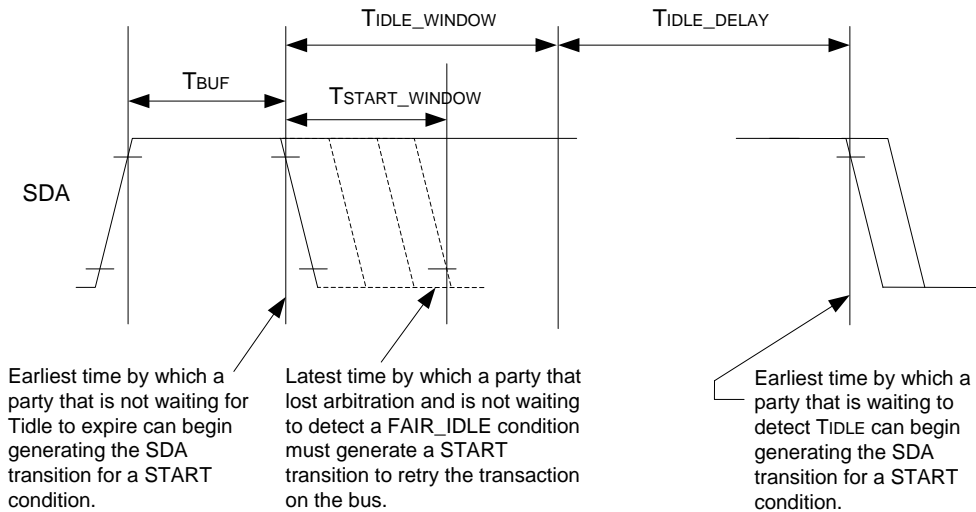
## 6.16 Fairness arbitration requirements for non-bridge endpoints

Non-bridge/bus owner endpoints ("simple endpoints") are required to implement the MCTP fairness arbitration extensions (when enabled) as follows:

- 606       • The endpoint's port shall support FAIR\_IDLE detection and implement the corresponding  
607       fairness policy.
  - 608       • Upon device power up or initialization, the endpoint does not need to detect a FAIR\_IDLE  
609       condition before first attempting to access the bus.
  - 610       • The endpoint cannot make the acceptance of a receive packet conditional on its ability to  
611       transmit a packet (that is, a simple endpoint shall not NACK incoming packets because it is  
612       trying to send an outgoing packet).
- 613       Meeting this requirement may require the endpoint to have separate transmit and receive  
614       buffers. This is the recommended implementation.
- 615       If a device is severely limited in buffer space and cannot allocate separate space for both  
616       transmit and received data, options are for the endpoint to allow its buffer to be over-written by  
617       the receive packet, or in some cases the endpoint may elect to do a dummy receive of the  
618       incoming packet (that is, ACK the incoming bytes, but internally drop them as they are coming  
619       in.)
- 620       • Higher layer protocols shall be used to handle the case when the endpoint is targeted by more  
621       messages than it can process. The buffering requirement for the MCTP Control Protocol  
622       messages is defined in the MCTP Base Specification. Buffering requirements for other message  
623       types are defined in the respective specifications for the message type.
  - 624       • An endpoint is allowed to NACK a packet if it is temporarily unable to accept it (for example,  
625       because of an *input* buffer-full condition). This should typically only occur if the endpoint is the  
626       target of packets from more than one source endpoint.
  - 627       • There is no direct limit of how long a non-bridge endpoint is allowed to successively NACK  
628       incoming packets. However, there are limits on how many packet-level retries a transmitter will  
629       attempt before it drops the transmitted packet, as well as message type-specific limits on how  
630       long and how many times a given message will be retried.
  - 631       • If an endpoint has an output transmit packet and it NACKs an input receive packet from lack of  
632       input buffer space, it shall be able to issue a START condition to transmit the output packet  
633       within  $T_{START\_WINDOW}$  after the bus becomes free, unless the endpoint is waiting to detect  $T_{IDLE}$ .
  - 634       • An endpoint that receives a NACK while it is performing a Master Write operation is not required  
635       to immediately conclude the Master Write operation and drop off the bus. The endpoint may  
636       continue the write operation through its conclusion. In either case, the master shall always  
637       conclude its transaction with a STOP condition, unless some other device on the bus first  
638       produces a START or STOP condition. The latter situation is an erroneous condition on the bus,  
639       but bridges shall be able to handle it. Devices shall always recognize START and STOP  
640       conditions regardless of the transaction or bit position on which they occur.
  - 641       • Endpoints that are NACK'd or lose arbitration shall retry transaction for PN1 retries (see  
642       Table 8).

643 **6.17 Fairness arbitration timing**

644 Figure 4, Table 5, and Table 6 present the specifications for the timing intervals for fairness arbitration on  
 645 SMBus and I<sup>2</sup>C relative to the data (SDA) signal. Refer to [SMBus](#) and [I<sup>2</sup>C](#) for the additional specifications  
 646 on the relationship between SCL and SDA for STOP, bus idle, and START conditions.



647

648 **Figure 4 – Fairness arbitration timing measurement for SMBus and I<sup>2</sup>C**

649 **Table 5 – Fairness arbitration timing values for 100 kHz SMBus/I<sup>2</sup>C**

Symbol	Min	Max	Unit	Notes
T <sub>BUF</sub>	4.7	–	µs	Per <a href="#">SMBus</a> 100 kHz specification
T <sub>START_WINDOW</sub>	–	20	µs	Window of time within which a device that is not waiting to detect a FAIR_IDLE condition shall generate START if the device is retrying to gain bus access after losing arbitration.
T <sub>IDLE_WINDOW</sub>	30	60	µs	Window of time within which a device that is waiting to detect a FAIR_IDLE condition shall not detect a bus busy condition. A FAIR_IDLE condition exists when bus busy is not detected within this interval.
T <sub>IDLE_DELAY</sub>	31	–	µs	A device that detects FAIR_IDLE condition shall wait this delay before attempting to generate START. This delay accommodates the difference between the T <sub>IDLE_WINDOW</sub> intervals implemented by different devices on the bus, plus additional time to accommodate bus skews between devices that are generating START and devices that are monitoring for it. This guarantees that one party that has detected T <sub>IDLE_WINDOW</sub> does not generate START before other devices that are detecting FAIR_IDLE have completed checking for their T <sub>IDLE</sub> window. Otherwise, the other devices would not see a FAIR_IDLE condition even though one occurred. (Therefore T <sub>IDLE_DELAY</sub> shall be greater than the difference between the T <sub>IDLE_WINDOW</sub> maximum and minimum.)



650

**Table 6 – Fairness arbitration timing values for 400 kHz I<sup>2</sup>C**

Symbol	Min	Max	Unit	Notes
T <sub>BUF</sub>	1.3	–	μs	Per I <sup>2</sup> C 400 kHz specification
T <sub>START_WINDOW</sub>	–	4	μs	Window of time within which a device that is not waiting to detect a FAIR_IDLE condition shall generate START if the device is retrying to gain bus access after losing arbitration.
T <sub>IDLE_WINDOW</sub>	5	20	μs	Window of time within which a device that is waiting to detect a FAIR_IDLE condition shall not detect a bus busy condition. A FAIR_IDLE condition exists when bus busy is not detected within this.
T <sub>IDLE_DELAY</sub>	16	–	μs	Device that detects FAIR_IDLE condition shall wait for this delay before attempting to generate START. This delay accommodates the difference between the T <sub>IDLE_WINDOW</sub> intervals implemented by different devices on the bus, plus additional time to accommodate bus skews between devices that are generating START and devices that are monitoring for it. This guarantees that one party that has detected T <sub>IDLE</sub> does not generate START before other devices that are detecting T <sub>IDLE</sub> have completed their T <sub>IDLE</sub> window. Otherwise, the other devices would not see a FAIR_IDLE condition even though one occurred. (Therefore T <sub>IDLE_DELAY</sub> shall be greater than the difference between the T <sub>IDLE_WINDOW</sub> maximum and minimum.)

651

**Table 7 – Fairness arbitration timing values for 1MHz I<sup>2</sup>C**

Symbol	Min	Max	Unit	Notes
T <sub>BUF</sub>	0.5	–	μs	Per I <sup>2</sup> C 1MHz specification
T <sub>START_WINDOW</sub>	–	2	μs	Window of time within which a device that is not waiting to detect a FAIR_IDLE condition shall generate START if the device is retrying to gain bus access after losing arbitration.
T <sub>IDLE_WINDOW</sub>	3	6	μs	Window of time within which a device that is waiting to detect a FAIR_IDLE condition shall not detect a bus busy condition. A FAIR_IDLE condition exists when bus busy is not detected within this interval.
T <sub>IDLE_DELAY</sub>	3.1	–	μs	Device that detects FAIR_IDLE condition shall wait this delay before attempting to generate START. This delay accommodates the difference between the T <sub>IDLE_WINDOW</sub> intervals implemented by different devices on the bus, plus additional time to accommodate bus skews between devices that are generating START and devices that are monitoring for it. This guarantees that one party that has detected T <sub>IDLE_WINDOW</sub> does not generate START before other devices that are detecting FAIR_IDLE have completed checking for their T <sub>IDLE</sub> window. Otherwise, the other devices would not see a FAIR_IDLE condition even though one occurred. (Therefore T <sub>IDLE_DELAY</sub> shall be greater than the difference between the T <sub>IDLE_WINDOW</sub> maximum and minimum.)

652 **6.18 MCTP packet timing requirements**

653 The timing specifications shown in Table 8 are specific to MCTP packet transfers on SMBus. Timing is  
 654 specified for a "point-to-point" connection. That is, timing is specified as if there were only two endpoints  
 655 in direct communication on the bus. In particular, the timing specifications assume that there is no clock  
 656 stretching that occurs due to other parties on the bus.

Table 8 – Timing specifications for MCTP packets on SMBus/I<sup>2</sup>C

Timing Specification	Symbol	Value	Description
Endpoint packet level retries	PN1	8	Number of times a non-bridge endpoint shall retry sending an MCTP packet upon receiving a NACK during the specified window (see Figure 3). An endpoint that gets successive NACKs shall do one retry for each NACK up to at least this number of retries. This also includes bridges when bridges are transmitting as an endpoint (as opposed to a bridge transmitting from its routing functionality).
Bridge packet level retries	PN2	12	Number of times an MCTP bridge (when transmitting packet for routing) shall retry sending an MCTP packet upon receiving a NACK during the specified window (see 6.14). A bridge shall do one retry on each NACK up to this number.
Packet transaction originator duration	PT1a	250 $\mu$ s per byte <sup>[1]</sup>	Overall duration shall be less than the specified interval times the number of bytes in the packet, starting from the byte following the slave byte through and including the PEC byte. Individual data byte transmissions may exceed the specification provided the cumulative duration for the packet is met.
Originator slave address byte duration	PT1b	250 $\mu$ s <sup>[1]</sup>	Amount of time, including any clock stretching, used to transmit the slave address, Wr, and ACK bits on the bus.
Slave-induced clock stretching	PT1c	250 $\mu$ s per byte <sup>[1]</sup>	MCTP devices that are receiving MCTP packets shall not clock stretch the overall packet more than the specified amount.  Note that MCTP devices may share the bus with non-MCTP SMBus devices that cause clock stretching that exceeds this specification.
The PT2 parameters are intended to help guide a controller in determining when it is acceptable to initiate a Master Write transaction if the controller powers up or initializes itself on a bus segment that may already be active. It also helps controllers know when it is acceptable to continue under conditions where a STOP condition may have been lost because a controller dropped off the bus due to an error condition. An implementation shall meet at least one of specifications PT2a or PT2b.			
Time-out waiting for bus free without seeing a STOP condition (Bus free determined by not detecting START or STOP)	PT2a	100 ms	For controllers that have hardware that can only detect bus-free/busy-busy status by monitoring for START and STOP conditions, the controller can assume the bus is free if PT2a seconds goes by without detecting a START or STOP condition.  If a START condition is detected, the time-out interval restarts.  If a STOP condition is detected, the controller can assuming the bus is free following the TBUF interval specified in <a href="#">SMBus</a> .  NOTE This interval effectively places an upper limit on the duration of a single transaction. The byte count in an MCTP packet limits the size of the transaction to 260 bytes. 100 ms is more than sufficient to cover this transfer.

Timing Specification	Symbol	Value	Description
Time-out waiting for bus free without seeing a STOP condition (Bus free determined by data/clock activity)	PT2b	50 $\mu$ s	The <a href="#">SMBus</a> specification defines a bus-free (idle) condition as TBUF seconds after a STOP condition, or by the data and clock lines being high for PT2b seconds (where the value for PT2b is taken from T <sub>HIGH, max</sub> as defined in <a href="#">SMBus</a> ).  If a controller has appropriate hardware support, monitoring PT2b and TBUF can be used to determine the bus-free (idle) condition in lieu of PT2a. This is generally the most efficient and highest performance way to detect bus free on SMBus.  SYSTEM IMPLEMENTATION NOTE: If "bit banded" I <sup>2</sup> C devices may be used on the same segment, it is important to ensure that those devices do not drive the clock and data high for more than T <sub>HIGH, max</sub> seconds during transactions.
SDA Low Timeout	PT3	2 sec min, 5 sec max	Time for a bus owner to monitor the SDA low level for a "Stuck 0" before attempting to clear the condition. (See 6.20.)
NOTE 1: Intervals include the ACK bit associated with the byte.			

658 **6.19 MCTP control message timing requirements**

659 The following timing specifications are specific to MCTP control messages on SMBus/I<sup>2</sup>C. Timing is  
 660 specified for a "point-to-point" connection. That is, timing is specified as if there were only two endpoints  
 661 in direct communication on the bus. In particular, the timing specifications assume that there is no clock  
 662 stretching occurs due to other parties on the bus.

663 Response specifications are given assuming that the requester is able to operate at full speed on the bus.  
 664 That is, clock stretching, if any, is solely generated by the requester.

665 Responses are not retried. A "try" or "retry" of a request is defined as a complete transmission of the  
 666 MCTP control message.

667 **Table 9 – Timing specifications for MCTP control messages on SMBus**

Timing Specification	Symbol	Min	Max	Description
Endpoint ID reclaim	T <sub>RECLAIM</sub>	5 sec	–	Minimum time that a bus owner shall wait before reclaiming the EID for a non-responsive hot-plug endpoint.
Number of request retries	MN1	2	See descr.	Total of three tries, minimum: the original try plus two retries. The maximum number of retries for a given request is limited by the requirement that all retries shall occur within MT4, max of the initial request.

Timing Specification	Symbol	Min	Max	Description
Request-to-response time	MT1	–	100 ms	This interval is measured at the responder from the end of the reception of the MCTP Control Protocol request to the beginning of the transmission of the response. This requirement is tested under the condition where the responder can successfully transmit the response on the first try.
Time-out waiting for a response	MT2	MT1 max+ 2*MT3 max	MT4, min <sup>[1]</sup>	This interval is measured at the requester from the end of the successful transmission of the MCTP Control Protocol request to the beginning of the reception of the corresponding MCTP Control Protocol response. This interval at the requester sets the minimum amount of time that a requester should wait before retrying an MCTP Control Protocol request.  Note: This specification does not preclude an implementation from adjusting the minimum time-out waiting for a response to a smaller number than MT2 based on measured response times from responders. The mechanism for doing so is outside the scope of this specification.
Transmission Delay	MT3	-	100 ms	Time to take into account transmission delay of an MCTP Control Protocol Message. Measured as the time between the end of the transmission of an MCTP Control Protocol message at the transmitter to the beginning of the reception of the MCTP Control Protocol message at the receiver.
Inter-Packet delay for Multi-Packet messages	MT3a		100ms	Allowed time measured from the end of the transmission of an MCTP packet with EOM=0 to the beginning of the following MCTP packet of the same Message (see Message assembly in <a href="#">DSP0236</a> ), measured at the transmitter
Instance ID expiration interval	MT4	5 sec <sup>[2]</sup>	6 sec	Interval after which the instance ID for a given response will expire and become reusable if a response has not been received for the request. This is also the maximum time that a responder tracks an instance ID for a given request from a given requester.

Timing Specification	Symbol	Min	Max	Description
NOTE 1: Unless otherwise specified, this timing applies to the mandatory and optional MCTP commands.				
NOTE 2: If a requester is reset, it may produce the same sequence number for a request as one that was previously issued. To guard against this, it is recommended that sequence number expiration be implemented. Any request from a given requester that is received more than MT4 seconds after a previous, matching request should be treated as a new request, not a retry.				

668

669 **6.20 "Stuck 0" condition handling**

670 A possible condition exists in SMBus and I<sup>2</sup>C where a slave device that is being read or is driving ACK  
 671 could be left driving a low (0) level onto the data line (SDA) of the bus. The bus uses a "wire OR'd"  
 672 approach, where the low (0) level takes precedence over the high (1) level. Therefore, if one party drives  
 673 a low (0) level onto the bus, the bus cannot go to a high (1) level until the low level is released.

674 This means that no other transactions can occur until this condition is cleared (because generating a  
 675 START or STOP condition on the bus requires being able to drive a high-to-low or low-to-high transition  
 676 on the data line, respectively).

677 This condition can occur due to the premature termination of a transaction from the master (as could  
 678 happen on device resets, power cycles, or firmware restarts, for example) or could occur due to the loss  
 679 of a clock due to electrical noise.

680 Effectively, what happens is that the device that was being accessed does not recognize that the  
 681 transaction has been terminated or that a clock was missed. The device continues to drive the 0 onto the  
 682 bus because it is waiting to get more clocks from the master to conclude the transaction, but those clocks  
 683 will never come unless some bus master takes steps to generate them.

684 The solution to this condition is to have a master clock the bus until the SDA line goes high, at which point  
 685 the master can issue a START or STOP condition to get the bus back in synchronization.

686 To accomplish this, the master needs to be able to access and clock the bus without paying attention to  
 687 the present state of the SDA line.

688 Many microcontrollers have the ability to have firmware dynamically reconfigure their SMBus pins as  
 689 general purpose I/O pins. If this is supported, it is straightforward for firmware to generate the necessary  
 690 clocks on the SCL line by bypassing the SMBus controller hardware and using programmed I/O to control  
 691 the pins instead. The firmware would then simply clock the bus until it sees a "1" condition on the SDA  
 692 line and then a new SMBus transaction can be launched.

693 NOTE It is recommended that MCTP bus owners include a provision to detect and clear Stuck 0 conditions on  
 694 SMBus buses that they own. The controller should do this if it can detect that a constant 0 condition has  
 695 existed on the SDA line for more than PT3 seconds.

696 **6.21 MCTP over SMBus/I<sup>2</sup>C protocol anti-aliasing**

697 MCTP over SMBus has been designed to allow one endpoint to support multiple protocols, such as ASF,  
 698 IPMI, or legacy device-specific protocols with a single slave address. The following clauses describe  
 699 provisions that can help support implement MCTP over SMBus in devices that also need to support other  
 700 SMBus or I<sup>2</sup>C protocols.

701 **6.21.1 IPMI**

702 The IPMI protocols for SMBus (IPMI over SMBus) and I<sup>2</sup>C (Intelligent Platform Management Bus, IPMB)  
 703 use the fourth byte of the transaction as a Source Slave Address byte, as does MCTP over SMBus.  
 704 However, the IPMI protocols require the least significant bit of that byte to be 0<sub>b</sub>, whereas MCTP over

705 SMBus requires the bit to be 1b. Thus, a device that needs to differentiate between MCTP over SMBus  
706 and the IPMI SMBus/I<sup>2</sup>C protocols can do so using that bit.

### 707 6.21.2 ASF

708 MCTP over SMBus uses the ASF specification reserved value of 0x0F for the command byte. Thus, the  
709 ASF-defined commands that use SMBus block-write protocol can be differentiated from MCTP over  
710 SMBus block-write using the command byte value. If necessary, other ASF SMBus write transactions,  
711 such as those for legacy sensor and control access can be differentiated from MCTP packets based on  
712 the length of the transaction. The ASF transactions are all shorter.

### 713 6.21.3 Integrating MCTP with legacy SMBus functions

714 This clause describes some possible options if MCTP is being added to a device that shall also support  
715 functions using a non-MCTP SMBus interface.

716 In general, there should be no problems having those functions co-exist with MCTP provided that the  
717 legacy SMBus operations do not require generating or accepting write transactions that use the MCTP  
718 value of 0x0F.

719 If the SMBus device currently uses the 0x0F MCTP command value for a device-specific purpose and it  
720 wants to use the same slave address, the following can be done:

- 721 • The device-specific command can be moved to a different command value. This is generally the  
722 most straightforward approach if it can be supported.
- 723 • Depending on the device-specific command definition, it may be possible to differentiate  
724 between the command and MCTP packets based on other differences, such as the overall  
725 length of the command or differences between the values in the fourth or fifth bytes of the  
726 command. (MCTP always uses 1b as the least significant bit of the fourth byte, and the fifth  
727 byte holds a fixed 4-bit value for the Header Version.)
- 728 • The device can implement MCTP over SMBus on a separate slave address from the legacy  
729 functions.

## 730 6.22 Well-known and reserved slave addresses

731 For bus segments that support ARP-able devices, Table 10 summarizes addresses that are generally  
732 reserved by SMBus or I<sup>2</sup>C and should either be avoided by devices. In addition, some are reserved (not  
733 to be used as a general device slave address) because those addresses are related to functions that are  
734 used by MCTP.

735 **Table 10 – Well-known and reserved slave addresses**

Slave Address bits [7:1]	R/W# bit [0]	Hex <sup>[7]</sup>	Comment	Disposition
0000 000	0	0x00	I <sup>2</sup> C general call address, IPMI broadcast	avoid <sup>[1]</sup>
0000 000	1	0x01	START byte	avoid <sup>[2]</sup>
0000 001	x	0x02 , 0x03	CBUS address	avoid <sup>[3]</sup>
0000 010	x	0x04 , 0x05	Address reserved for different bus format	avoid
0000 011	x	0x06 , 0x07	Reserved for future use by I <sup>2</sup> C specifications	avoid

Slave Address bits [7:1]	R/W# bit [0]	Hex <sup>[7]</sup>	Comment	Disposition
0000 1XX	X	0x08–0x0F	I <sup>2</sup> C specification, high-speed mode master code	avoid <sup>[4]</sup>
0001 000	X	0x10	SMBus host	rsvd
0001 100	X	0x18, 0x19	SMBus Alert Response address	rsvd
0010 000	X	0x20, 0x21	IPMI BMC address	avoid <sup>[5]</sup>
0101 000	X	0x50, 0x51	Reserved for ACCESS.bus host	avoid (ACCESS.bus defunct)
0110 111	X	0x6E, 0x6F	Reserved for ACCESS.bus default address	avoid
1111 0XX	X	0xF0–0xF7	I <sup>2</sup> C 10-bit slave addressing <sup>[1]</sup>	avoid <sup>[6]</sup>
1111 1XX	X	0xF8–0xFF	Reserved for future use by I <sup>2</sup> C specifications	avoid
1100 001	X	0xC2, 0xC3	SMBus Device Default address	rsvd. Used for SMBus ARP with MCTP
<p><b>NOTE 1.</b> This address is used as a broadcast address in IPMI and I<sup>2</sup>C. It should be avoided if IPMI management controllers may be used on the same bus segment. In I<sup>2</sup>C, it is reserved for two purposes: to broadcast a portion of an address that is used for devices that have a portion of their address that is configurable, and as an optional mechanism for a device to master and broadcast its slave address onto the bus. MCTP does not support the use of this address for the I<sup>2</sup>C address assignment or slave address broadcast purposes.</p> <p><b>NOTE 2.</b> The I<sup>2</sup>C START byte is a pre-amble to the slave address that is intended to provide time for firmware driven I<sup>2</sup>C interfaces to shift into polling of I<sup>2</sup>C clock and data lines after a START condition has been detected. This is a very rarely used option in I<sup>2</sup>C. MCTP does not support the use of the START byte with MCTP or non-MCTP devices.</p> <p><b>NOTE 3.</b> CBUS is an ancestor of I<sup>2</sup>C, developed by Philips Semiconductor. It uses a data and clock signal similar to I<sup>2</sup>C, but with a third signal (SEN) used to generate the START and STOP conditions on the bus. This address range was reserved by the I<sup>2</sup>C specification to enable a degree of backward compatibility with CBUS devices sharing the I<sup>2</sup>C SCL and SDA lines as the CBUS clock and data lines, respectively. While listed as a reserved address in the I<sup>2</sup>C specification, few SMBus/I<sup>2</sup>C implementations using MCTP will have any need to also support CBUS devices.</p> <p><b>NOTE 4.</b> MCTP is not defined to support I<sup>2</sup>C high-speed mode operation.</p> <p><b>NOTE 5.</b> This address is the "well known address" for an IPMI BMC. This address should be avoided if an IPMI BMC may be used on the same MCTP segment.</p> <p><b>NOTE 6.</b> Used in conjunction with the R/W# bit position to deliver the most-significant three address bits for I<sup>2</sup>C 10-bit addressing. MCTP protocols and data structures do not support 10-bit addressing on SMBus or I<sup>2</sup>C segments. MCTP only supports 7-bit addresses for MCTP and non-MCTP devices on a bus segment.</p> <p><b>NOTE 7.</b> By convention, when the 7-bit slave address field is represented as a two-digit hexadecimal number, it is treated as an 8-bit value where the 7-bit address occupies the upper 7 bits and the least significant bit is 0b or 1b according to the value of the SMBus/I<sup>2</sup>C Read-Write bit associated with the slave address.</p>				

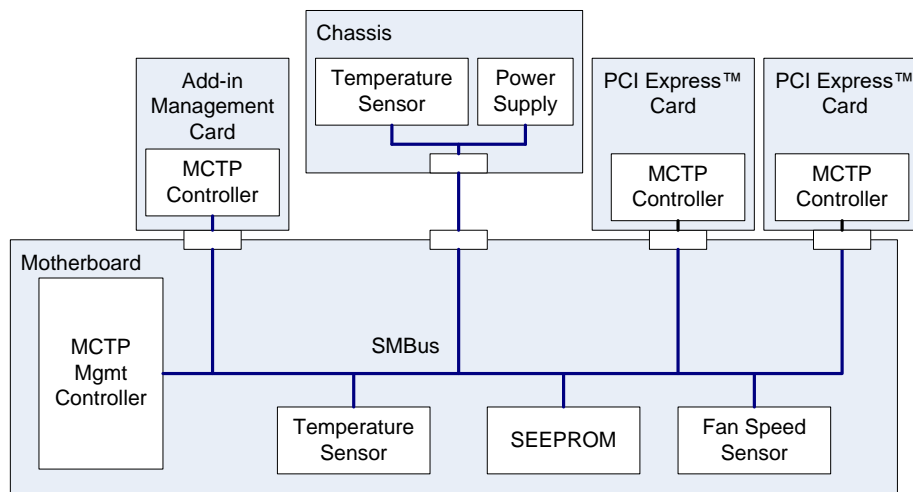
736 **6.23 Fixed address allocation**

737 One of the problems that an implementer often faces is choosing which slave address to use. For the  
 738 PCI™ and PCI Express™ bus specifications, the specifications require that devices on standard  
 739 connectors defined by those specifications have their addresses set through SMBus ARP. Therefore,  
 740 fixed address allocation is not an option for PCI add-in cards themselves. In fixed bus implementations,  
 741 however, there are many situations where it is desired or necessary to utilize fixed-address devices.

742 From a practical point-of-view, SMBus and I<sup>2</sup>C do not have an effective central registry or other  
 743 mechanism for avoiding conflicts in the assignment and use of slave addresses among device vendors.

744 While there are potential registries of device slave address usage for SMBus (under the System  
 745 Management Interface Forum) and I<sup>2</sup>C (from Philips Semiconductor), these have not generally been used  
 746 by device vendors and there is no group or standard that works to enforce conformance to those  
 747 registries.

748 Most device vendors provide a configurable range of three or more addresses to enable an implementer  
 749 to reconcile address conflicts on a single segment. Because typically only a small number of fixed-  
 750 address devices are used on a given segment, it is frequently possible to configure devices so they do  
 751 not have overlapping addresses. This approach is problematic, however, in situations where a component  
 752 that is attached to that segment in the platform may come from several sources. Clause 6.23 provides  
 753 guidelines to allocating fixed addresses that are designed to reduce the number of conflicts that could  
 754 occur when multiple suppliers provide different elements of a computer system (see Figure 5 for an  
 755 example).



756

757

**Figure 5 – Example system configuration**

## 758 6.24 Recommended address range allocation for computer systems

759 This clause provides a recommended allocation of SMBus addresses between board, chassis, and add-in  
 760 uses that help avoid address conflicts when fixed addresses are used. It also serves as a general  
 761 guideline of what addresses a generic ARP master should use for allocation to PCI/PCIe add-in cards.

762 There might be cases when MCTP is used within a typical computer system application where the  
 763 motherboard may come from one supplier, the chassis from another supplier, and possibly add-in  
 764 modules from yet another supplier. To facilitate the mix-and-match of these elements and to help avoid  
 765 the need for every system manufacturer to set up their own address allocation conventions with suppliers,  
 766 MCTP recommends that system manufacturers follow the address allocation approach initially defined by  
 767 the IPMI specifications (see Table 11). This approach splits the available fixed addresses (addresses  
 768 other than reserved addresses) into four main usage areas:

769 **B Board:** An area reserved for board set manufacturer use (where *board set* would be the  
 770 motherboard and other boards that accompany that motherboard from the same vendor).



771       **C**       **Chassis:** An area reserved for use by vendors that make chassis in which a third-party board  
772                   set would be used.

773       **A**       **Add-in:** For third-party add-in devices (for example, modules or add-in cards that used fixed  
774                   addresses and would be used in combination with a motherboard or chassis where there is a  
775                   connection to a SMBus segment implementing MCTP).

776                   NOTE    PCI/PCIe add-in cards that use standard PCI connectors are required to support SMBus ARP  
777                   and fixed addresses are not used.

778       **R**       **Reserved** for IPMI, I<sup>2</sup>C, SMBus, or MCTP uses. Includes the *avoid* addresses from Table 10.

779    By following this convention, future motherboards can offer connections to chassis elements and third-  
780    party modules where those devices can use fixed addresses, if required. It also provides a convention to  
781    avoid conflicts if legacy non-MCTP devices share the same SMBus segment.

Table 11 – Slave address allocation for computer systems

Use	Address	Typical Device	Use	Address	Typical Device
R	0x00	I <sup>2</sup> C, IPMB broadcast	C		
	0x01	I <sup>2</sup> C			
	0x02	I <sup>2</sup> C		0x48	SMBUS/I2C IO Expander, such as 8574
	0x04-0x0E	I <sup>2</sup> C		0x4A	SMBUS/I2C IO Expander, such as 8574
	0x20	IPMB uC (BMC)		0x4C	SMBUS/I2C IO Expander, such as 8574
	0x50	ACCESS.bus		0x4E	SMBUS/I2C IO Expander, such as 8574
	0x6E	ACCESS.bus		0x52-0x6C	58h, 5Ah, 5Ch = Heceta
	0xF0-0xF6	I <sup>2</sup> C			
	0xF8-0xFE	I <sup>2</sup> C			
A	0x10	SMBus host (B)	0x78	SMBUS/I2C IO Expander, such as 8574A	
	0x12-0x16		0x7A	SMBUS/I2C IO Expander, such as 8574A	
	0x18	SMBus Alert Response address (B)	0x7Ch	SMBUS/I2C IO Expander, such as 8574A	
	0x1A-0x1E		0x7E	SMBUS/I2C IO Expander, such as 8574A	
	0x30-0x3E		0x9A	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621	
			0x9C	uC (pri. HSC), DS1624, DS1621	
	0xD0-0xDE		0x9E	uC (sec. HSC), DS1624, DS1621	
B	0x22	uC (FPC, ICMB) <sup>[1]</sup>	0xA0-0xA2	FRU (Power Supply FRU or SEEPROM)	
	0x24	uC (PBC) <sup>[1]</sup>	0xAC	SEEPROMSEEPROM	
	0x26		0xAE	SEEPROM	
	0x28	SM Card <sup>[1]</sup>	0xB0-0xB2	Power Supply Device (PMBus)	
	0x2A-0x2E		0xE8-0xEE	I2C Bus Switch	
	0x40	SMBUS/I2C IO Expander, such as 8574A	NOTE 1: Term from IPMI usage. FPC = front panel controller, PBC = Power Backplane Controller, ICMB = Intelligent Chassis Management Bus bridge, SM Card = System Management Card		
	0x42	SMBUS/I2C IO Expander, such as 8574A			
	0x44	SMBUS/I2C IO Expander, such as 8574A			
	0x46	SMBUS/I2C IO Expander, such as 8574A			
	0x70	SMBUS/I2C IO Expander, such as 8574A			
0x72	SMBUS/I2C IO Expander, such as 8574A				

Use	Address	Typical Device	Use	Address	Typical Device
	0x74	SMBUS/I2C IO Expander, such as 8574A			
	0x76	SMBUS/I2C IO Expander, such as 8574A			
	0x80-0x8E				
	0x90	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621, 8591			
	0x92	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621, 8591			
	0x94	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621, 8591			
	0x96	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621, 8591			
	0x98	TEMPERATURE SENSORS, SUCH AS LM75, DS1624, DS1621, 8591			
	0xA4	EEPROM			
	0xA6	EEPROM			
	0xA8	EEPROM			
	0xAA	EEPROM			
	0xC0				
	0xC2	<a href="#">SMBus</a> Device Default address			
	0xC4-0xCE				
	0xE0-0xE6	I2C Bus Switch			

## ANNEX A (informative)

### Notation

783  
784  
785  
786  
787

#### 788 Notations

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

- 790 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets  
791 starting from byte two and continuing to and including byte N. The lowest offset is on  
792 the left, the highest is on the right.
- 793 • (6) Parentheses around a single number can be used in message field descriptions to  
794 indicate a byte field that may be present or absent.
- 795 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range  
796 may be present or absent. The lowest offset is on the left, the highest is on the right.
- 797 • PCIe Underlined, blue text is typically used to indicate a reference to a document or  
798 specification called out in 2, "Normative References" or to items hyperlinked within the  
799 document.
- 800 • rsvd Abbreviation for "reserved." Case insensitive.
- 801 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets  
802 are given as zero-based values (that is, the least significant bit [LSb] offset = 0).
- 803 • [7:5] A range of bit offsets. The most significant bit is on the left, the least significant bit is  
804 on the right.
- 805 • 1b The lower case "b" following a number consisting of 0s and 1s is used to indicate the  
806 number is being given in binary format.
- 807 • 0x12A A leading "0x" is used to indicate a number given in hexadecimal format.

808

809  
810  
811  
812  
813

## ANNEX B (informative)

### Change log

Version	Date	Description
1.0.0	2009-07-28	
1.1.0	2017-02-16	Added 1MHz speed mode Moved NACK window to a separate clause Corrected timing parameters description Updated reserved addresses mapping in Table 11 Updated Table 2 Updated revisions for the normative references
1.2.0	2020-04-06	Added timing parameter MT3a Removed redundant definition of ARP term Updated hyperlink to SMBus spec

814