

12 Copyright notice

13 Copyright © 2017 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29 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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 PCI-SIG, PCIe, and the PCI HOT PLUG design mark are registered trademarks or service marks of PCI-
33 SIG.

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

35 For information about patents held by third-parties which have notified the DMTF that, in their opinion,
36 such patent may relate to or impact implementations of DMTF standards, visit
37 <http://www.dmtf.org/about/policies/disclosures.php>.

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

39

40 CONTENTS

41 Foreword 5

42 Introduction..... 6

43 Document conventions..... 6

44 Typographical conventions 6

45 ABNF usage conventions 6

46 1 Scope 7

47 2 Normative references 7

48 3 Terms and definitions 8

49 4 Symbols and abbreviated terms..... 8

50 5 Conventions 9

51 5.1 Reserved and unassigned values..... 9

52 5.2 Byte ordering..... 9

53 6 NC-SI over MCTP overview 9

54 6.1 NC-SI over RBT 10

55 6.2 NC-SI over MCTP 11

56 7 NC-SI over MCTP specific considerations 12

57 7.1 Packages and channels..... 12

58 7.2 Routing of NC-SI Pass-through traffic 13

59 7.2.1 Transmit NC-SI Pass-through traffic (MC to LAN)..... 13

60 7.2.2 Receive NC-SI Pass-through traffic (LAN to MC)..... 13

61 7.3 Multiple NC support 13

62 7.3.1 NC-SI arbitration support 13

63 7.4 Flow control..... 14

64 7.4.1 Flow control for MCTP packets..... 14

65 7.4.2 Flow control for NC-SI over MCTP Control messages 14

66 7.4.3 Flow control for NC-SI Pass-through packets. 14

67 7.5 Interleaving of messages 14

68 7.5.1 Interleaving of MCTP Control and NC-SI messages 14

69 7.5.2 Interleaving of NC-SI Control and Ethernet over MCTP messages 14

70 7.6 Ordering rules for NC to MC traffic 15

71 7.7 Assembly requirements 16

72 7.8 Multiple MCTP transport bindings..... 16

73 7.8.1 Supported message types over different MCTP transport bindings 16

74 7.8.2 MCTP EID and physical address changes. 17

75 7.8.3 NC discovery flows 17

76 7.8.4 MC update flow 18

77 7.8.5 Transition between mediums 18

78 7.9 Package selection 20

79 8 Supported NC-SI commands 21

80 9 Message types 24

81 9.1 NC-SI message type (0x02) 24

82 9.1.1 Encapsulation 24

83 9.1.2 Version 25

84 9.2 Ethernet message type (0x03)..... 25

85 9.2.1 Encapsulation 25

86 9.2.2 Version 26

87 10 NC-SI support specific to MCTP transport 27

88 10.1 Get Supported Media command (0x54)..... 27

89 10.2 Get Supported Media response (0xD4) 27

90 10.3 Transport Specific AENs Enable (0x55) 28

91 10.4 Transport Specific AENs Enable response (0xD5)..... 29

92 10.5 Medium change AEN 29

93 11 Packet-Based Timing Specific to MCTP Binding 30

94 ANNEX A (informative) Notation and conventions 31

95 ANNEX B (informative) Change log 32

96

97 **Figures**

98 Figure 1 – NC-SI over RBT traffic flow diagram 10

99 Figure 2 – NC-SI over MCTP traffic flow diagram 11

100 Figure 3 – Single MCTP EID to multiple NC-SI channels mapping 12

101 Figure 4 – Multiple MCTP EIDs to multiple NC-SI channels mapping 13

102 Figure 5 – Multiple MCTP transport bindings example 16

103

104 **Tables**

105 Table 1 – MCTP Message types for NC-SI over MCTP 11

106 Table 2 – Supported NC-SI commands 21

107 Table 3 – NC-SI messages encapsulation 24

108 Table 4 - MCTP Transport Header fields 25

109 Table 5 – MCTP Specific Message Header field 25

110 Table 6 – Ethernet messages encapsulation 26

111 Table 7 - MCTP Transport Header fields 26

112 Table 8 – MCTP Specific Message Header field 26

113 Table 9 – Get Supported Media command packet format 27

114 Table 10 – Get Supported Media response packet format 27

115 Table 11 – Get Supported Media response media descriptors format 28

116 Table 12 – Transport Specific AENs Enable command packet format 28

117 Table 13 – Transport Specific AENs enable field format 29

118 Table 14 – Transport Specific AENs Enable response packet format 29

119 Table 15 – Medium change AEN format 30

120 Table 16 – NC-SI timing parameters specific to MCTP Binding 30

121

Foreword

122 The *NC-SI over MCTP Binding Specification* (DSP0261) was prepared by the Platform Management
123 Components Intercommunications (PMCI Working Group) of the DMTF.

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

126 **Acknowledgments**

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

128 **Editor:**

- 129 • Eliel Louzoun – Intel Corporation

130 **Contributors:**

- 131 • Alan Berenbaum – SMSC
- 132 • Patrick Caporale - Lenovo
- 133 • Philip Chidester – Dell Inc.
- 134 • Kelly Couch – Intel Corporation
- 135 • Benzi Friedman – Intel Corporation
- 136 • Yuval Itkin – Mellanox Technologies
- 137 • Ed Klodnicki – IBM
- 138 • Koe Kozlowski – Dell Inc.
- 139 • Patrick Kutch – Intel Corporation
- 140 • John Leung – Intel Corporation
- 141 • Patrick Schoeller – Hewlett Packard Enterprise
- 142 • Hemal Shah – Broadcom Limited
- 143 • Tom Slaight – Intel Corporation
- 144 • Bob Stevens – Dell Inc.

145

146

Introduction

147 The *NC-SI over MCTP Binding Specification* defines new MCTP messages used to convey NC-SI Control
148 packets and Ethernet traffic over MCTP to allow NC-SI Pass-through traffic over MCTP. This specification
149 is based on the [DSP0222](#) specification and uses the same NC-SI Control packet definitions.

150 Document conventions

151 Typographical conventions

152 The following typographical conventions are used in this document:

- 153 • Document titles are marked in *italics*.
- 154 • Important terms that are used for the first time are marked in *italics*.
- 155 • Terms include a link to the term definition in the "Terms and definitions" clause, enabling easy
156 navigation to the term definition.
- 157 • ABNF rules are in `monospaced font`.

158 ABNF usage conventions

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

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

163

164

NC-SI over MCTP Binding Specification

1 Scope

166 The *NC-SI over MCTP Binding Specification* defines the bindings between NC-SI protocol elements and
167 MCTP elements in order for NC-SI Control and Pass-Through traffic to be transported using MCTP.

168 Portions of this specification rely on information and definitions from other specifications, which are
169 identified in clause 2. Two of these references are particularly relevant:

- 170 • DMTF [DSP0222](#), *Network Controller Sideband Interface (NC-SI) Specification*, provides the
171 NC-SI base control that is to be bound over MCTP by this specification.
- 172 • DMTF [DSP0236](#), *Management Component Transport Protocol (MCTP) Base Specification*,
173 defines the MCTP transport on which the NC-SI Control and Pass-through packets are to be
174 conveyed.

2 Normative references

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

180 Unless otherwise specified, for DMTF documents this means any document version that has minor or
181 update version numbers that are later than those for the referenced document. The major version
182 numbers must match the major version number given for the referenced document.

183 DMTF DSP0222, *Network Controller Sideband Interface (NC-SI) Specification 1.1*
184 http://www.dmtf.org/sites/default/files/standards/documents/DSP0222_1.1.pdf

185 DMTF DSP0223, *Generic Operations 1.0*,
186 http://www.dmtf.org/standards/published_documents/DSP0223_1.0.pdf

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

189 DMTF DSP0237, *Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding*
190 *Specification 1.1*
191 http://www.dmtf.org/sites/default/files/standards/documents/DSP0237_1.1.pdf

192 DMTF DSP0238, *Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding*
193 *Specification 1.0*
194 http://www.dmtf.org/standards/published_documents/DSP0238_1.0.pdf

195 DMTF DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.4*
196 http://www.dmtf.org/standards/published_documents/DSP0239_1.4.pdf

197 DMTF DSP1001, *Management Profile Specification Usage Guide 1.2*,
198 http://www.dmtf.org/standards/published_documents/DSP1001_1.2.pdf

199 ACPI, *Advanced Configuration and Power Interface Specification Revision 4.0a*, April 5, 2010
200 <http://www.acpi.info/DOWNLOADS/ACPIspec40a.pdf>

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

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

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

207 **3 Terms and definitions**

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

210 The terms "shall" ("required"), "shall not," "should" ("recommended"), "should not" ("not recommended"),
211 "may," "need not" ("not required"), "can" and "cannot" in this document are to be interpreted as described
212 in [ISO/IEC Directives, Part 2](#), Annex H. The terms in parenthesis are alternatives for the preceding term,
213 for use in exceptional cases when the preceding term cannot be used for linguistic reasons. Note that
214 [ISO/IEC Directives, Part 2](#), Annex H specifies additional alternatives. Occurrences of such additional
215 alternatives shall be interpreted in their normal English meaning.

216 The terms "clause," "subclause," "paragraph," and "annex" in this document are to be interpreted as
217 described in [ISO/IEC Directives, Part 2](#), Clause 5.

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

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

222 Refer to [DSP0222](#) for terms and definitions that are used in the NC-SI specification.

223 The terms defined in [DSP0004](#), [DSP0223](#), and [DSP1001](#) apply to this document. The following additional
224 terms are used in this document.

225 **3.1 Endpoint**

226 An MCTP endpoint unless otherwise specified.

227 **3.2**

228 **System Power States**

229 **S0 and Sx**

230 S0 represents an active system

231 Sx represents system power states S1 – S5, which reflects various levels of inactivity of a system.

232 The definition of the power states is as defined in [ACPI](#).

233 **4 Symbols and abbreviated terms**

234 The abbreviations defined in [DSP0004](#), [DSP0223](#), and [DSP1001](#) apply to this document. The following
235 additional abbreviations are used in this document.

236 **4.1**

237 **ACPI**

238 Advanced Configuration and Power Interface

239	4.2
240	IANA
241	Internet Assigned Numbers Authority
242	4.3
243	FCS
244	Frame Check Sequence
245	4.4
246	MCTP
247	Management Component Transport Protocol
248	4.5
249	MC
250	Management Controller
251	4.6
252	NC
253	Network Controller
254	4.7
255	NC-SI
256	Network Controller Sideband Interface
257	4.8
258	RID
259	PCIe Requester ID (Bus/Device/Function).

260 **5 Conventions**

261 **5.1 Reserved and unassigned values**

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

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

266 **5.2 Byte ordering**

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

269 **6 NC-SI over MCTP overview**

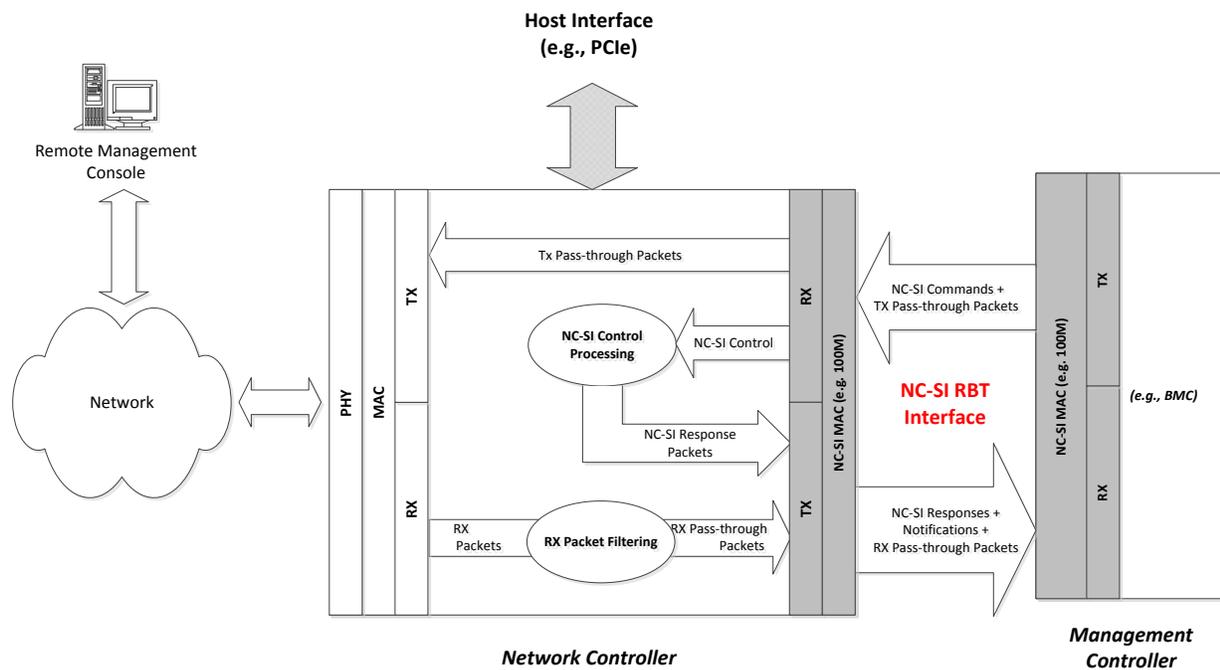
270 NC-SI over MCTP is based on DSP0222 ([NC-SI](#)). The *NC-SI over MCTP Binding Specification* replaces
271 the RBT Protocol with a definition of NC-SI communications using MCTP. The MCTP Transport Bindings
272 are defined in other companion specifications such as *MCTP SMBus Binding Specification* ([DSP0237](#))
273 and *MCTP PCIe Binding Specification* ([DSP0238](#)). Only the NC-SI command processing is inherited from
274 DSP0222. Thus only parts of the [NC-SI](#) specification not related to the physical transport protocol may be
275 relevant to this specification.

276 **6.1 NC-SI over RBT**

277 A Network Controller Sideband Interface (NC-SI) is a combination of logical and physical paths that
 278 interconnect the Management Controller and Network Controller(s) for the purpose of transferring
 279 management communication traffic among them. NC-SI includes commands and associated responses,
 280 which the Management Controller uses to control the status and operation of the Network Controller(s).
 281 NC-SI also includes a mechanism for transporting management traffic and asynchronous notifications.

282 Figure 1 depicts the NC-SI Traffic Flow Diagram as currently defined by [NC-SI](#). As indicated, the interface
 283 is based on RBT. The figure depicts a single management controller and a single Ethernet device, which
 284 contains a single port. [NC-SI](#) comprehends multiple Network Controller devices (or “packages”) and ports
 285 (or “channels”).

286



287

288 **Figure 1 – NC-SI over RBT traffic flow diagram**

289

290 The DSP0222 specification can be divided in three parts. The first two parts are defined as RBT:

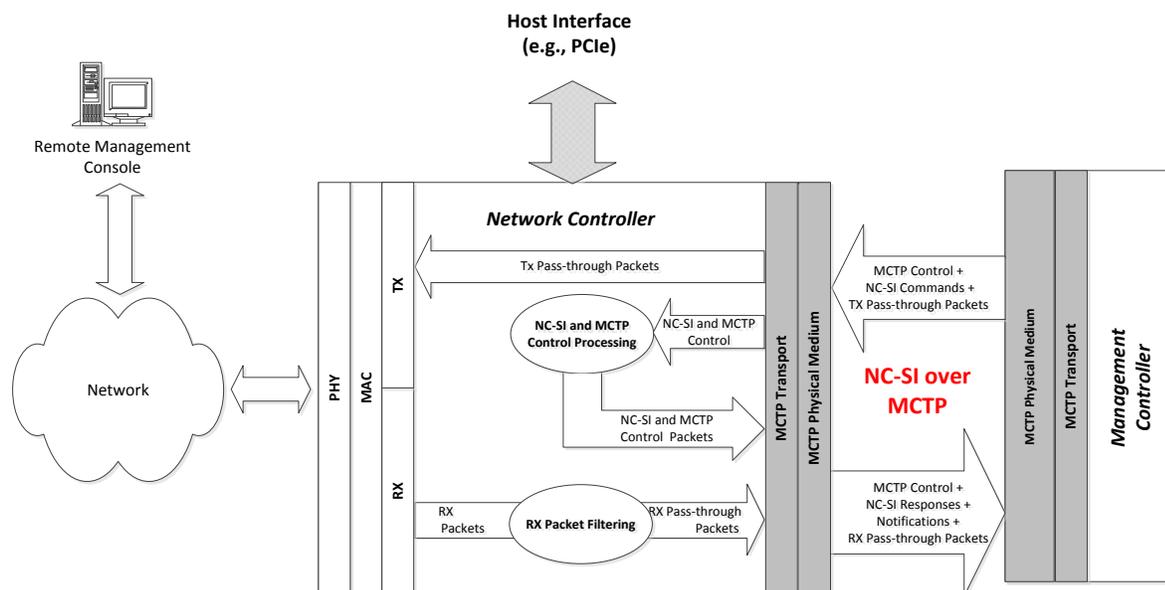
- 291
- A physical layer based on enhancements to the RMI specification.
 - 292 • A transport layer based on Ethernet packets. This layer allows differentiation of control frames
 293 based on a specific Ethertype (0x88F8).
 - 294 • A control protocol defining a set of commands allowing an MC to configure and monitor Network
 295 Controllers and their Pass-through channels for MC to network communication. The command
 296 set functionality can be extended using OEM commands.

297 **6.2 NC-SI over MCTP**

298 NC-SI over MCTP replaces the transport layer defined in NC-SI with MCTP. The physical layer used is
 299 one of the transport bindings on which MCTP can be bound (for example, PCIe or SMBus).

300 Figure 2 shows a possible architecture that provides equivalent functionality to [\[NC-SI\]](#) over MCTP. The
 301 NC-SI MAC block in each device is replaced by an MCTP block and a Medium-specific block. The MCTP
 302 block handles MCTP messages. The Medium-specific blocks consist of whatever layers are involved in
 303 mapping MCTP to an underlying medium such as SMBus, PCIe, or USB. Because the layering for each
 304 medium may be unique in its constitution and terminology, a generic single block is depicted.

305



306
 307 **Figure 2 – NC-SI over MCTP traffic flow diagram**

308
 309 The differentiation between NC-SI Control and Pass-through packets is achieved by using two different
 310 MCTP message types as defined in [DSP0239](#) and listed in Table 1.

311 **Table 1 – MCTP Message types for NC-SI over MCTP**

Message Type	Message Type Code	Description
NC-SI Control	0x02	Messages used to encapsulate NC-SI Control traffic (commands, responses, and AEN) over MCTP
Ethernet	0x03	Messages used to encapsulate Ethernet traffic (for example, NC-SI Pass-through) over MCTP

312
 313 Both NC-SI Control and Ethernet types of MCTP messages can be conveyed over multiple MCTP
 314 packets.

315 The encapsulation of NC-SI Control traffic in MCTP messages is described in subclause 9.1.1. The
 316 encapsulation of NC-SI Pass-through traffic in MCTP messages is described in subclause 9.2.1.

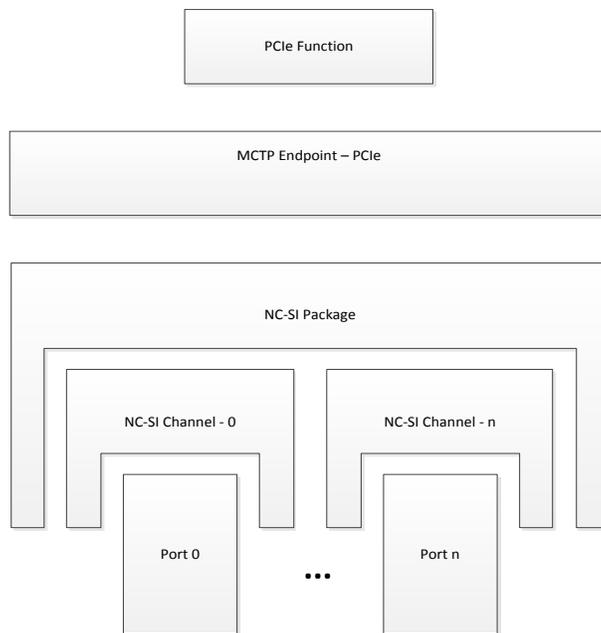
317 7 NC-SI over MCTP specific considerations

318 7.1 Packages and channels

319 The NC-SI specification defines different topologies using the concepts of channels and packages. A
 320 channel is associated with a network port and a package is usually associated with a physical device that
 321 exposes a single NC-SI bus. In an MCTP context, a package is related to an MCTP endpoint. Typically, a
 322 package is identified by a single MCTP EID on an MCTP network.

323 **NOTE** Each device may expose multiple MCTP endpoints on different transport bindings (for example PCIe and
 324 SMBus). The EID on each transport binding may be different. In this case, the NC-SI package may be associated
 325 with multiple EIDs but only a single EID shall be used for NC-SI over MCTP at a given moment.

326 For example, each MCTP endpoint is associated with a PCIe endpoint and its physical address (as
 327 defined in [DSP0238](#)) in an MCTP over PCIe VDM transport binding implementation. A multi-function PCIe
 328 device has multiple physical addresses available. Such a device may choose to expose one NC-SI
 329 package with multiple NC-SI channels via a single MCTP PCIe endpoint (as described in Figure 3) or
 330 multiple NC-SI packages, each package with a single NC-SI channel exposed via a dedicated MCTP
 331 PCIe endpoint (as described in Figure 4).



332

333 **Figure 3 – Single MCTP EID to multiple NC-SI channels mapping**

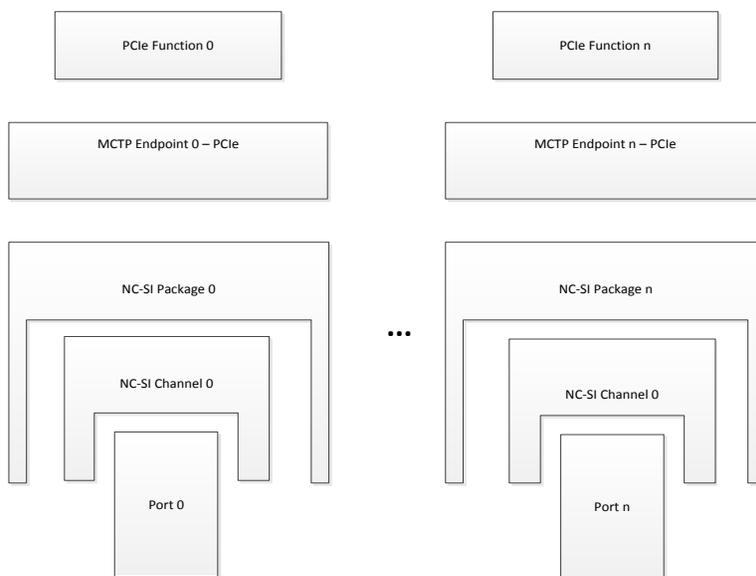


Figure 4 – Multiple MCTP EIDs to multiple NC-SI channels mapping

334

335

336

337 Multiple MCTP transport bindings handling is described in subclause 7.8.

338 NOTE All the MCTP message segmentation and reassembly capabilities required are defined at the package level.

339 **7.2 Routing of NC-SI Pass-through traffic**

340 **7.2.1 Transmit NC-SI Pass-through traffic (MC to LAN)**

341 Because multiple NC-SI channels can share an EID, identification of channel is still based on the source
 342 MAC address of the packet. Given the shared media behavior of RBT in multidrop configurations, packets
 343 not destined to this package can be seen. In NC-SI over MCTP, the NC-SI pass-through packets are
 344 routed over an MCTP network, thus packets destined to other packages are not expected. The NC should
 345 drop the received NC-SI TX Pass-through packets that are not destined to its package and may count
 346 them in one of the channels' Tx error counter. If counted, these errors shall be included in the "Pass-
 347 through TX Packets Dropped" counter as part of the Get NC-SI Pass-through Statistics Response.

348 **7.2.2 Receive NC-SI Pass-through traffic (LAN to MC)**

349 The forwarding of network traffic to the MC shall use the same rules as defined in DSP0222.

350 **7.3 Multiple NC support**

351 **7.3.1 NC-SI arbitration support**

352 In the original NC-SI specification, hardware and command-based arbitration are defined as ways to
 353 share an inherently point-to-point media between different NCs. With MCTP, the media itself may provide
 354 other means to arbitrate between different NCs. Thus, there is no need to use NC-SI HW arbitration
 355 method to arbitrate between multiple NCs on an MCTP network.

356 An NC supporting the NC-SI over MCTP binding shall retain the support for the 'select package' and
 357 'deselect package' commands to allow control of asynchronous transmission from the NC.

358 7.4 Flow control

359 7.4.1 Flow control for MCTP packets

360 A physical medium supporting NC-SI over MCTP communication shall be able to buffer at least one NC-
361 SI Control or Ethernet message at the rate of the physical layer. Flow control of MCTP packets between
362 the Network Controller and the Management Controller (if any) may be handled by the flow control
363 mechanisms that are specified for that particular MCTP Transport Binding for a physical medium. For
364 example, a network controller may use the SMBus clock stretching mechanism to delay the reception of
365 MCTP packets or may drop such packets.

366 7.4.2 Flow control for NC-SI over MCTP Control messages

367 Flow control of NC-SI Control over MCTP messages is handled by the request/response protocol used for
368 those messages. The Network Controller shall be able to process a single NC-SI command at a time from
369 the Management Controller. The Management Controller shall wait until getting a NC-SI response to that
370 NC-SI command, or for a response timeout, before sending another NC-SI command over MCTP to that
371 NC.

372 7.4.3 Flow control for NC-SI Pass-through packets.

373 The NC-SI Pass-through traffic flow control used in RBT is an Ethernet-specific technology that is not well
374 suited to an MCTP transport. An implementation of this specification may support Ethernet flow control,
375 but it will apply only to Ethernet messages (message type – 0x3) and not to messages of NC-SI Control
376 over MCTP type (message type – 0x2). The method used to control the rate of transmission of Ethernet
377 packets is beyond the scope of this specification.

378 7.5 Interleaving of messages

379 7.5.1 Interleaving of MCTP Control and NC-SI messages

380 According to the MCTP specification [[MCTP](#)], an endpoint shall accept MCTP Control messages that are
381 interleaved among NC-SI Control over MCTP or Ethernet over MCTP message packets. This is to avoid
382 scenarios where functions such as the MCTP bus owner are 'locked out' from managing the MCTP bus
383 because of NC-SI Pass-through traffic.

384 Correspondingly, MCTP Control Message responses shall be able to be interleaved among incoming NC-
385 SI Control over MCTP or Ethernet over MCTP message packet. However, the MCTP Control Message
386 responses may be held up and transmitted between Ethernet Messages, provided that the MCTP
387 command request-to-response timing requirements are met.

388 7.5.2 Interleaving of NC-SI Control and Ethernet over MCTP messages

389 NC-SI Control over MCTP and Ethernet over MCTP messages to the same EID shall not be interleaved.
390 Similar to the [DSP0222](#) specification case, NC-SI Control and Ethernet packets are interleaved at the
391 message level. An MC operating with multiple NC may interleave messages sent to different NCs.
392

393 **7.6 Ordering rules for NC to MC traffic**

394 The following table defines which type of messages should pass other types of packets to avoid
 395 deadlocks. The decisions are done at a message level. Interleaving within messages is defined in the
 396 previous sections. The following behaviors are expected:

- 397 • Yes—the second message (row) shall be allowed to pass the first (column) to avoid deadlock
 398 (When blocking occurs, the second message is required to pass the first message)
- 399 • Y/N—there are no requirements. The second message may optionally pass the first message or
 400 be blocked by it as long as the timing specifications for the messages are met.
- 401 • No—the second message shall not be allowed to pass the first message. This is required to avoid
 402 out of order events.

403

Row Pass Column?	MCTP Control Response (Col 1)	NC-SI Response (Col 2)	NC-SI AEN (Col 3)	Ethernet Packet (Col 4)
MCTP Control Response (Row A)	Y/N	Y/N	Yes	Y/N
NC-SI response Packet (Row B)	Y/N	Y/N	Yes	Y/N
NC-SI AEN (Row C)	Y/N	Y/N	No	Y/N
Ethernet Packet (Row D)	Y/N	Y/N	Y/N	No

404 **Notes** (The letter and number indicates the row and column in the table above):

405

- 406 • **A** This row relates only to the precedence of MCTP base control messages
 407 over NC-SI and Ethernet messages and not over other MCTP message
 408 types.
- 409 • **A1** This situation will occur only in NCs accepting multiple outstanding
 410 MCTP control commands.
- 411 • **B2** This situation will occur only in NCs accepting multiple outstanding NC-
 412 SI commands.
- 413 • **A3, B3** An NC-SI AEN may be blocked if the channel is disabled or the package
 414 deselected. Thus it should not block MCTP Control or NC-SI responses.
- 415 • **C3** AENs should be sent in order of occurrence to avoid cases where the
 416 latest received status is obsolete. For example in the case of a link-down
 417 event followed by a link-up event, the AEN on the link-up event must not
 418 pass the AEN on the link-down event.
- 419 • **D4** Ethernet packets must be sent in order to avoid out-of-order events in
 420 the upper layers.

421 7.7 Assembly requirements

422 According to the interleaving requirements described in subclause 7.5, the NC shall be able to assemble
 423 a single NC-SI Control or Ethernet over MCTP message at a time. The maximum Ethernet packet size is
 424 defined in subclause 9.2. The maximum NC-SI packet size is defined in subclause 9.1.

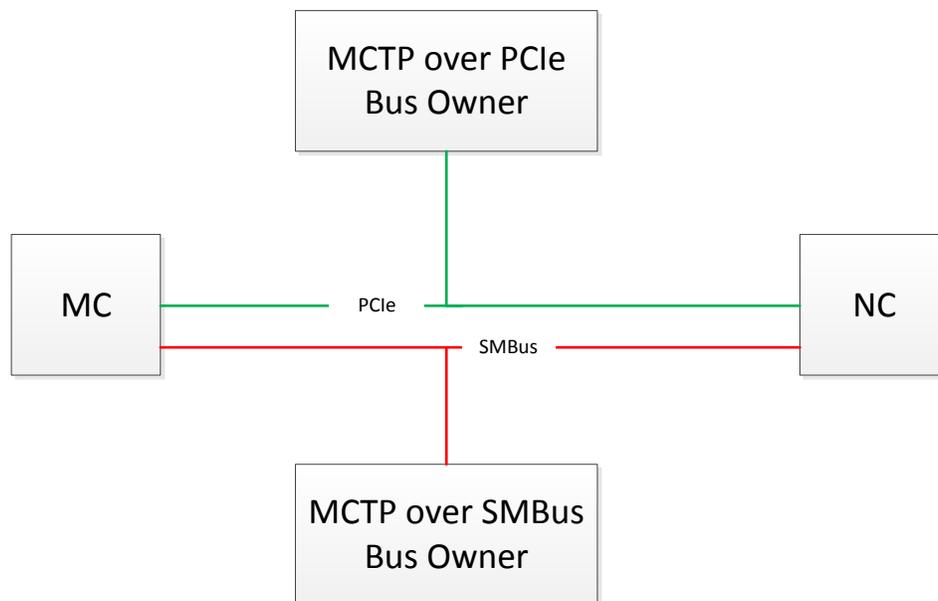
425 Buffering requirements for other message types are not covered in this specification.

426 7.8 Multiple MCTP transport bindings

427 In the DSP0222 specification, the channels use a single physical interface all the time. In NC-SI over
 428 MCTP, multiple MCTP transport bindings may be used at different times to convey NC-SI traffic to allow
 429 tradeoffs between data rate and power consumption. The following requirements apply to those MCTP
 430 transport bindings:

- 431 1) NC-SI control messages (identified by MCTP message type 0x2) shall be supported
- 432 2) Ethernet messages (identified by MCTP message type 0x3) may be supported

433 Figure 5 shows an example of multiple MCTP transport bindings using MCTP over PCIe VDM and MCTP
 434 over SMBus. The types of NC-SI over MCTP traffic on each MCTP transport binding may vary as
 435 described above.



436

437 **Figure 5 – Multiple MCTP transport bindings example**

438 7.8.1 Supported message types over different MCTP transport bindings

439 An endpoint may support different MCTP message types over different MCTP transport bindings. For
 440 example, an NC may choose to support Ethernet message type over MCTP PCIe VDM transport only. It
 441 is recommended that an MC initially determines the supported message types on a given medium during
 442 the discovery phase using the Get Supported Message Type MCTP Control command prior to
 443 transmitting MCTP traffic of specific MCTP message type on the medium.

444

445 **7.8.2 MCTP EID and physical address changes.**

446 The NC-SI package mapping of the NC or the MC to MCTP EID and/or physical interface address may
447 change due to the following reasons:

- 448 1) Changes in the MCTP transport medium used. For example moving from PCIe to SMBus
449 medium when PCIe becomes unavailable for MCTP communication due to change of
450 power state.
- 451 2) Changes in the EID to physical address mapping. For example when changing medium or
452 during re-enumeration process or in a multifunction PCIe device, if the function of which
453 RID is being used is disabled by the host, the MCTP endpoint may move to another
454 function.

455 In order to avoid breakup of network connections, and in order to avoid the need to reconfigure the NC,
456 the NC-SI connection should be kept alive during the transition. The MC is responsible for the
457 reconnection of the channel in case of address mapping changes. The next clause describes possible
458 flows that may be used to ease the re-discovery of an NC whose address has changed. A flow by which
459 the MC can expose a change of its own address to the NC(s) is described in subclause 7.8.4.

460 According to the [MCTP](#) specification, an MC or NC that has its physical address changed should send an
461 MCTP Discovery Notify command to the bus owner so that the routing tables can be updated.

462 **7.8.3 NC discovery flows**

463 The MC may use one of the following example flows to discover a NC whose address has changed.

464 **7.8.3.1 Full discovery**

465 The simplest and most time consuming method is to discover the NC partner by using the standard
466 MCTP discovery method. This method works with NCs that support at least MCTP 1.1 and NC-SI 1.0.

467 The following flow may be used:

- 468 • The MC detects a potential address update condition (for example: power state change, link
469 status change, or re-enumeration) or detects an NC-SI timeout condition (as defined in section
470 6.8.2.1 of [NC-SI](#)).
- 471 • The MC finds all the endpoints in the system by sending an MCTP “Get Routing Table Entries”
472 command to the bus owner and to any bridges in the MCTP network.
- 473 • For each device listed, the MC checks whether it supports the required MCTP message types
474 (NC-SI Control and optionally Ethernet) by using the MCTP “Get Message Type Support”
475 command.

476 For each potential endpoint discovered by using the method above, the MC checks whether it is the
477 original NC partner, for example by sending an “Get Version ID” NC-SI command to the original NC ID
478 and checking the response.

479 **7.8.3.2 UUID based discovery**

480 This method is based on the usage of the “Resolve UUID” MCTP command.

481 To use this method, the bus owner or bridge must support the “Resolve UUID” MCTP command and the
482 NC must support the “Get Endpoint UUID” MCTP command.

483 The following flow may be used:

- 484 • When the NC-SI channel is first established by using some proprietary method (for example by
485 using the flow from the previous section), the MC may send a “Get Endpoint UUID” MCTP
486 command to the NC. It then keeps the UUID information for future use.
- 487 • MC periodically sends a “Get Routing Table” Command to the bus owner to receive updated
488 endpoints addresses.
- 489 • The NC whose address changes or that wants to move to another active bus sends a “Discovery
490 Notify” MCTP command to the bus owner of the new bus.
- 491 • As part of the routing table update, the bus owner sends a “Get Endpoint UUID” MCTP command
492 to the NC and updates its routing table accordingly.
- 493 • The MC sends a “Resolve UUID” MCTP Command to the bus owner by using the previously
494 saved NC UUID. In response, it gets the list of EIDs matching this UUID.
- 495 • The MC can check if the relevant message types (NC-SI Control and optionally Ethernet) are
496 supported on the new bus by using an MCTP “Get Message Type Support” command.
- 497 • The MC may then send any NC-SI Command to the NC to communicate with the NC on the new
498 medium.

499 7.8.3.3 NC-SI based discovery

500 The NC must support the “Get Supported Media” NC-SI command as defined in clause 10.1 to use this
501 method.

502 The following flow may be used.

- 503 • The MC detects a potential address update condition (for example: power state change, link
504 status change, AEN from the NC, or re-enumeration) or detects a timeout condition on NC-SI (as
505 defined in section 6.8.2.1 of [NC-SI](#)).
- 506 • If the original bus is still available (for example, when transitioning from SMBus to PCIe), it may
507 send on the original bus a “Get Supported Media” NC-SI command. In the response, the NC will
508 provide information on the routing that should be used on the new bus and on the support for
509 Pass-through on this bus.
- 510 • The MC may then send any NC-SI Command to the NC to communicate with the NC on the new
511 medium.

512 This method may not be applicable when there is no active channel that can be used to send the “Get
513 Supported Media” NC-SI command over. In this case, one of the other methods should be used.

514 7.8.4 MC update flow

515 In the case where MC physical address or its MCTP EID changes, it may send an “Enable Channel” NC-
516 SI command to the NC. This command will update the MC EID and physical address used by the NC.

517 7.8.5 Transition between mediums

518 A transition of an NC-SI package from one medium to another can occur due to changes in the available
519 media. For example, a transition from SMBus to PCIe can occur when PCIe becomes available to provide
520 a larger bandwidth.

- 521 A transition of an NC-SI from one medium to another is achieved when the NC is deselected on the first
522 medium and selected on the second medium as described in subclause 7.9.
- 523 The NC may notify the MC about the state of a medium using an AEN.
- 524 1) Potential loss of a medium prior to losing the medium
- 525 2) Availability of a new medium
- 526 Alternatively, the MC may be aware of the medium change independently, for example, by detecting its
527 own PCIe bus became active, by interaction with the BIOS, and so on.
- 528 The MC may initiate the transition by using MCTP Control and NC-SI Control messages as described in
529 subclause 7.8.5.1.
- 530 A transition can be between mediums with different levels of support of Ethernet MCTP messages.
- 531 When an NC transitions from a medium on which Ethernet messages were supported to a second
532 medium on which Ethernet messages are not supported, the NC should stop sending and receiving
533 Ethernet messages on the first medium after the NC-SI channel had been deselected on the first medium.
- 534 The MC may transition back later to the first medium for communicating Ethernet messages. If the MC
535 transitions back to the first medium supporting Ethernet messages, it may resume communications of
536 Ethernet messages based on the previous configuration. If the configuration was lost during the
537 transitions, the NC shall return to the NC-SI Initial State (as described in section 6.2.4 of [NC-SI](#)).
- 538 Even if NC-SI Pass-through traffic (Ethernet messages) is supported over multiple mediums, Pass-
539 through traffic shall not be transitioned to a new medium before the connection between the MC and the
540 NC is re-established on the new medium. The NC shall support the following flows to initiate a transition
541 to the new medium:
- 542 • If the current medium is still active (for example when moving from SMBus to PCIe to achieve
543 better throughput), the NC shall keep its Pass-through traffic on the original medium (both MC to
544 network and network to MC). The NC shall also send outstanding NC-SI responses on the
545 original medium.
546 NOTE The MC can stop the traffic from the NC on the current medium by sending “Disable Channel” and
547 “Disable Channel Network TX” NC-SI commands to all the channels before the transition. In this case, it
548 shall send “Enable Channel” and “Enable Channel Network TX” NC-SI commands to all active channels on
549 the new medium, to allow the traffic to resume.
 - 550 • If the current medium is inactive (for example, when moving from PCIe to SMBus due to a power
551 transition), the NC shall stop transferring Ethernet messages. If a packet is being transmitted by
552 the NC when the original medium becomes unavailable, the NC shall not continue the
553 transmission of the packet and the packet may be lost. Outstanding NC-SI responses may be
554 discarded by the NC.
 - 555 • When any NC-SI command is received from the MC on the new medium (apart from “*Deselect*
556 *Package*”), the NC shall move to “Selected” state on the new medium (see subclause 7.9).
 - 557 ○ An NC-SI Rx Pass-through message to the MC on the current medium shall be
558 completed by the NC on the current medium and only after that shall the NC send the
559 NC-SI response to the MC on the new medium.
 - 560 ○ The next Pass-through message sent to the MC after a successful response to the NC-SI
561 command shall be sent on the new medium.

562 • The NC shall accept Pass-through traffic from the MC on the new medium after the NC moves to
563 "Selected" state on the new medium and sends the first successful NC-SI response.

564 • The same algorithm as described above shall be used for the selection of the medium to use for
565 sending NC-SI AEN messages to the MC.

566 An NC that uses multiple MCTP transport bindings should support at least one of the UUID based
567 recovery or the NC-SI based recovery methods in addition to the Full Discovery mechanism.

568 **7.8.5.1 MC transition flow**

569 The MC can initiate a transition between mediums for one of the following reasons.

570 1) Loss of medium for NC-SI over MCTP communications. For example, system transitioning
571 into a low power state may make PCIe medium unavailable for NC-SI over MCTP
572 communications over PCIe VDM transport.

573 2) Reception of an AEN from the NC notifying a medium state change. For example, an NC
574 may notify the MC about the potential loss of the PCIe medium, triggering a transition to
575 SMBus.

576 The following flow can be used by the MC to initiate a transition between mediums:

577 • If the current medium is still active (for example when moving from SMBus to PCIe to achieve
578 better throughput), the MC can keep its traffic on the original medium until it discovers the NC by
579 using one of the flows described in subclause 7.8.3. If the current medium is inactive (for
580 example, when moving from PCIe to SMBus due to a power transition), the MC will stop
581 transferring Ethernet messages with NC until discovery of the NC.

582 • The MC can then send an "Enable Channel" NC-SI Command, or any other command to the NC
583 to select it on the new medium. The MC will then wait for the NC response before starting to send
584 packets on the new medium. The MC will complete transmission of the current Ethernet message
585 before sending the command and will not send Ethernet messages while waiting for the
586 response. The MC will accept Ethernet message on the original medium until the response from
587 the NC is received on the current medium.

588 • If Pass-through is supported by the NC over only a single medium, when transitioning out of this
589 medium, the MC will not send Pass-through traffic to the NC and will not expect to receive traffic
590 from the NC.

591 • If a medium becomes unavailable while an MC waits for an NC-SI command response, it can
592 assume the command was lost and retry it on the new medium.

593 **7.9 Package selection**

594 The "Selected" state of an NC-SI package is defined for each of the MCTP transports to which it can bind.
595 A package can be selected only on a single MCTP medium at a given point of time.

596 As in DSP0222, a package is selected by reception of a "Select Package" on the MCTP medium or any
597 other command except "Deselect Package".

598 A package is deselected on a specific MCTP medium by reception of a "Deselect Package" command,
599 selection of the package on another medium or if the physical medium on which it operate becomes
600 unavailable. If the packet is deselected by an NC-SI command it should move to the deselected state only
601 after sending a response to the command.

602 A package is allowed to send Ethernet messages or NC-SI Control messages on an MCTP medium only
603 if in the "Selected" state on that medium.

604 An NC should use the source EID and source physical address received from the last received NC-SI
 605 command to respond to this command and as the destination of subsequent Ethernet messages. If a
 606 command is received during the transmission of an Ethernet message, the destination should change
 607 only at the beginning of the next message.

608 The channel selection state and all other NC-SI configurations may be kept during the transition from one
 609 medium to another. If the configuration is altered during the transition, the NC shall return to Initial State.

610 8 Supported NC-SI commands

611 The supported NC-SI commands when bound to MCTP is a subset of the commands in [DSP0222](#)
 612 specification. The subset of supported commands varies according to the supported messages as
 613 indicated in the response to the Get Message Type Support MCTP Control command. If only the NC-SI
 614 Control message type is supported, the commands related to the Pass-through traffic control are not
 615 supported. If both the NC-SI Control and Ethernet message types are supported, these commands are
 616 supported. Table 2 lists the supported commands according to the supported message types.

617 Optional commands may have different implementation over different media.

618 Note that some commands are not applicable for MCTP binding and are listed here only for
 619 completeness. These commands are marked as “Not part of binding”.

620

621

Table 2 – Supported NC-SI commands

Command Type	Command Name	Description	Response Type	Command Support Requirement NC-SI Control Messages Only	Command Support Requirement NC-SI Control and Ethernet Messages
0x00	Clear Initial State	Used by the Management Controller to acknowledge that the Network Controller is in the Initial State	0x80	M	M
0x01	Select Package	Used to explicitly select a controller package to transmit packets through the NC-SI interface	0x81	O3	O ³
0x02	Deselect Package	Used to explicitly instruct the controller package to stop transmitting packets through the NC-SI interface	0x82	O ³	O ³
0x03	Enable Channel	Used to enable the NC-SI channel and to start the forwarding of bidirectional Management Controller packets	0x83	M	M
0x04	Disable Channel	Used to disable the NC-SI channel and to stop the forwarding of bidirectional Management Controller packets	0x84	M	M
0x05	Reset Channel	Used to synchronously put the Network Controller back to the Initial State	0x85	M	M

Command Type	Command Name	Description	Response Type	Command Support Requirement NC-SI Control Messages Only	Command Support Requirement NC-SI Control and Ethernet Messages
0x06	Enable Channel Network TX	Used to explicitly enable the channel to transmit Pass-through packets onto the network	0x86	N/A	M
0x07	Disable Channel Network TX	Used to explicitly disable the channel from transmitting Pass-through packets onto the network	0x87	N/A	M
0x08	AEN Enable	Used to control generating AENs	0x88	C	C
0x09	Set Link	Used during OS absence to force link settings, or to return to auto-negotiation mode	0x89	O	M
0x0A	Get Link Status	Used to get current link status information	0x8A	O	M
0x0B	Set VLAN Filter	Used to program VLAN IDs for VLAN filtering	0x8B	N/A	M
0x0C	Enable VLAN	Used to enable VLAN filtering of Management Controller RX packets	0x8C	N/A	M
0x0D	Disable VLAN	Used to disable VLAN filtering	0x8D	N/A	M
0x0E	Set MAC Address	Used to configure and enable unicast and multicast MAC address filters	0x8E	N/A	M
0x10	Enable Broadcast Filtering	Used to enable full or selective broadcast packet filtering	0x90	N/A	M
0x11	Disable Broadcast Filtering	Used to disable all broadcast packet filtering, and to enable the forwarding of broadcast packets	0x91	N/A	M
0x12	Enable Global Multicast Filtering	Used to disable forwarding of all multicast packets to the Management Controller	0x92	N/A	C
0x13	Disable Global Multicast Filtering	Used to enable forwarding of all multicast packets to the Management Controller	0x93	N/A	C
0x14	Set NC-SI Flow Control	Used to configure IEEE 802.3 flow control on NC-SI	0x94	N/A	O
0x15	Get Version ID	Used to get controller-related version information	0x95	M	M

Command Type	Command Name	Description	Response Type	Command Support Requirement NC-SI Control Messages Only	Command Support Requirement NC-SI Control and Ethernet Messages
0x16	Get Capabilities	Used to get optional functions supported by the NC	0x96	M ¹	M
0x17	Get Parameters	Used to get configuration parameter values currently in effect on the controller	0x97	M2	M
0x18	Get Controller Packet Statistics	Used to get current packet statistics for the Network Controller	0x98	O	O
0x19	Get NC-SI Statistics	Used to request the packet statistics specific to the NC-SI interface	0x99	O	O
0x1A	Get NC-SI Pass-through Statistics	Used to request NC-SI Pass-through packet statistics	0x9A	N/A	O
0x1B	Get Package Status	Used to get current status of the package	0x9B	O	O
0x50	OEM Command	Used to request vendor-specific data	0xD0	O	O
0x52	Get Package UUID	Returns a universally unique identifier (UUID) for the package	0xD2	O	O
0x53	PLDM	Used for PLDM request over NC-SI over RBT	0xD3	Not part of binding	Not part of binding
0x54	Get Supported Media	Used to return the media on which NC-SI can run and routing information for each medium.	0xD4	O	O
0x55	Transport Specific AEN Enable	Used to control generating Transport specific AENs	0xD5	O	O
<p>Key: M = Mandatory (required) O = Optional C = Conditional (see command description) N/A = Not applicable</p>					
<ol style="list-style-type: none"> The only part of the response that is relevant is the AEN control support field. The only part of the response that is relevant is the Link Settings, AEN control fields and the Channel Enabled flag in the Configuration Flags. The 'Select Package' and 'Deselect Package' commands impact only transmission of NC-SI Control and Ethernet over MCTP message types and do not impact other MCTP message types. 					

622 **9 Message types**

623 The MC and the NC shall support the NC-SI over MCTP message type (0x02). The Ethernet message
624 type should be supported if NC-SI Pass-through traffic is expected on this interface.

625 **9.1 NC-SI message type (0x02)**

626 This message type is used to carry NC-SI Control packets that are identified by the NC-SI Ethertype in
627 the DSP0222 specification. This includes command, response, and AEN packets.

628 The maximum NC-SI message payload size is 1500 bytes to keep the same limit as in [NC-SI](#). This
629 includes the payload starting from the MC ID field.

630 **9.1.1 Encapsulation**

631 The encapsulation of NC-SI Control packets includes the packet as described in the Control packet data
632 structure of [DSP0222](#) specification encapsulated in an MCTP header. NC-SI messages may be
633 fragmented to multiple MCTP packets.

634 NC-SI control packets communicated over MCTP do not follow the Ethernet frame encapsulation defined
635 in DSP0222 for NC-SI over RMIIBased Transport (RBT) transport binding. NC-SI control packets over
636 MCTP shall not include Ethernet frame header, Ethernet packet pad, and Ethernet Frame Check
637 Sequence (FCS). Instead, the encapsulation described in Table 3 shall be used to encapsulate NC-SI
638 control messages.

639 NOTE The Control packets frames in DSP0222 uses a DA, SA, and Ethertype MAC header. The DA and SA part
640 do not contain any useful data and the Ethertype is used to differentiate between Control packets and Ethernet traffic.
641 In NC-SI over MCTP, this Ethernet framing is not used, as the differentiation is achieved through usage of different
642 message types.

643

644 **Table 3 –NC-SI messages encapsulation**

Bytes	+0								+1								+2								+3							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
00..03	RSVD				Header Version				Destination Endpoint ID								Source Endpoint ID								S	E	Pkt		T	Message		
																								M	M	seq #	O	Tag				
04..07	IC	Message Type							MC ID								Header Revision								Reserved							
	0	0x02																														
08..11	IID								Command								Channel ID								Reserved				Payload Len			
12..15	Payload Length																Reserved															
16..19	Reserved																															
20..23	Reserved																Control Packet Payload															
...	...																															
...	...																Control Packet Payload								Payload Padding (as required)							
...	Payload Padding (as reqrd)																Checksum 3:1															
...	Checksum 0																															

645

646 See [NC-SI](#) for details of the NC-SI Control packets format.

647 The following tables describe the value for the various fields of the message whose description differs
648 from the description in the MCTP or NC-SI specification.

649

Table 4 - MCTP Transport Header fields

Field Name	Field Size	Value	Comment
Tag Owner (TO)	1 bit	Varies	Indicates that the Tag field value was generated by the message source = 0b Tag not from message source = 1b Tag from message source. Should be set for Commands and AEN packets. Should be cleared for Response packets.
Msg Tag	3 bits	Varies	The Tag field shall be set by the source of the message.

650

651

Table 5 – MCTP Specific Message Header field

Field Name	Field Size	Value	Comment
IC Bit	1 bit	0b	NC-SI over MCTP does not define message integrity check as it relies on the MCTP packet integrity check provided by the underlying medium or checks that are encapsulated in the message payload. This field is present only in the first packet of a message (SOM = 1).
Message Type	7 bits	0x02	Identifies the MCTP message type as an NC-SI Control over MCTP message. This field is present only in the first packet of a message (SOM = 1).

652 **9.1.2 Version**

653 The version that shall be reported for this message type in the Get MCTP Version Support response is
654 0xF1F0F000 (1.0.0)

655 **9.2 Ethernet message type (0x03)**

656 This message type is used to carry NC-SI Pass-through packets. Ethernet messages may be fragmented
657 to multiple MCTP packets.

658 The nominal Ethernet message size that shall be supported is 1518 bytes to accommodate a full Ethernet
659 packet including a VLAN but without FCS. If additional L2 tags are expected (for example, MACSec), the
660 supported packet size shall increase accordingly.

661 **9.2.1 Encapsulation**

662 The encapsulation of Ethernet packets includes the entire Ethernet frame from the Source MAC address
663 to the end of the payload, not including the FCS, prefixed with an MCTP header.

664 NOTE In [NC-SI](#), the FCS was required as part of the Ethernet encapsulation used over RMII. When Ethernet
665 packets are sent over other mediums, the medium specific error recovery mechanisms are used and the FCS is not
666 required.

667 The FCS should be added by the NC for packets sent by the MC to the network and should be checked
668 and removed by the NC for packets received from the network to the MC. Packets with a wrong FCS
669 should not be forwarded to the MC.

670 This behavior is consistent with the FCS offload provided by NCs to the host OS.

671 **Table 6 – Ethernet messages encapsulation**

Bytes	+0								+1								+2								+3							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
00..03	RSVD				Header Version				Destination Endpoint ID								Source Endpoint ID								SOM	EOM	Pkt seq #	TO	Message Tag			
04..07	IC 0	Message Type 0x03							Destination Address 5:3																							
08..11	Destination Address 2:0																Source Address 5															
12..15	Source Address 4:1																															
16..	Source Address 0								Optional L2 tags																							
...	Optional L2 tags								Ethertype								Ethernet Payload															
...	Ethernet Payload (no FCS)																															

672

673 The following tables describe the value for the various fields of the message whose description differs
 674 from the description in the MCTP or NC-SI specification.

675 **Table 7 - MCTP Transport Header fields**

Field Name	Field Size	Value	Comment
Tag Owner (TO)	1 bit	1b	Indicates that the Tag field value was generated by the message source = 0b Tag not from message source = 1b Tag from message source Should be set for all packets
Msg Tag	3 bits	Varies	The Tag field shall be set by the source of the message.

676 **Table 8 – MCTP Specific Message Header field**

Field Name	Field Size	Value	Comment
IC Bit	1 bit	0b	NC-SI over MCTP does not define a message integrity check because it relies on the MCTP packet integrity check provided by the underlying medium or checks that are encapsulated in the message payload. This field is present only in the first packet of a message (SOM = 1).
Message Type	7 bits	0x03	Identifies the MCTP message type as an Ethernet over MCTP message. This field is present only in the first packet of a message (SOM = 1).

677 **9.2.2 Version**

678 The version that shall be reported for this message type in the Get MCTP Version Support response is
 679 0xF1F0F000 (1.0.0)

680 **10 NC-SI support specific to MCTP transport**

681 The following commands and AEN may be implemented as part of this specification to allow an
 682 implementation of the discovery flow described in clause 7.8.3.3.

683 **10.1 Get Supported Media command (0x54)**

684 This command is used to query a device about the Media on which NC-SI can be conveyed. This
 685 command is optional and is applicable only if more than one media is supported.

686 The Get Supported Media command is addressed to the package, rather than to a particular channel (that
 687 is, the command is sent with a Channel ID where the Package ID subfield matches the ID of the intended
 688 package and the Internal Channel ID subfield is set to 0x1F).

689 Table 9 illustrates the packet format of the Get Supported Media command.

690 **Table 9 – Get Supported Media command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Checksum			
20..45	Pad			

691 **10.2 Get Supported Media response (0xD4)**

692 In the absence of any error, the package shall process and respond to the Get Supported Media
 693 command by sending the response packet and payload shown in Table 10.

694 **Table 10 – Get Supported Media response packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved			Number of medias supported
20..	Media descriptors as described in Table 11. The number of media descriptors is according to the Number of medias supported field value.			
...	Checksum			
...	Pad			

695

696

Table 11 – Get Supported Media response media descriptors format

Byte	Description
0	EID. Should be 0x0 if Physical Medium Identifier is RBT.
1	Physical Transport Binding Identifier, according to MCTP ID specification (DSP0239). Should be 0x0 if Physical Medium Type Identifier is RBT.
2	Physical Medium Identifier, according to MCTP ID specification (DSP0239). This value is used to indicate what format the following physical address data is given in.
3.0	NC-SI Pass-through is supported. 0: NC-SI Pass-through is not supported over this medium. 1: NC-SI Pass-through is supported over this medium.
3.6:1	Reserved
3.7	Status 0: Medium is not available. 1: Medium is available.
4	Physical Address Size. Should be 0x0 if Physical Medium Identifier is RBT.
5:N	Physical Address. This field is not present if Physical Medium Identifier is RBT. If present, this field is valid only if the Status bit is set and its value is unspecified otherwise.

697 **10.3 Transport Specific AENs Enable (0x55)**

698 Network Controller implementations shall support this command on the condition that the Network
 699 Controller generates one or more transport specific AENs defined in this specification. The AEN Enable
 700 command enables and disables the different transport specific AENs supported by the Network
 701 Controller. The Network Controller shall copy the AEN MC ID field from the Transport Specific AEN
 702 Enable command into the MC ID field in every subsequent AEN sent to the Management Controller as
 703 defined in [DSP0222](#).

704 Table 12 illustrates the packet format of the Enable Transport Specific AENs command.

705 The current version of this command only supports the Medium Change AEN.

706

Table 12 –Transport Specific AENs Enable command packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved	AEN MC ID	Transport Specific AENs enable	
20..23	Checksum			
24..45	Pad			

707 **Table 13 –Transport Specific AENs enable field format**

Bit Position	Field Name	Value Description
0	Medium Change AEN Control (0x70)	0b = Disable Medium Change AEN 1b = Enable Medium Change AEN
1..15	Reserved For future AEN	Reserved

708 **10.4 Transport Specific AENs Enable response (0xD5)**

709 In the absence of any error, the package shall process and respond to the Transport Specific AENs
710 Enable command by sending the response packet and payload shown in Table 14.

711 **Table 14 –Transport Specific AENs Enable response packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Checksum			
...	Pad			

712

713 **10.5 Medium change AEN**

714 The Medium change AEN is used to alert the MC that there was a status change in one of the media
715 supported by the NC, or such a change is expected according to some external or internal condition
716 detected by the NC.

717 This AEN should be sent if any change occurred in the status of one of the media supported by the
718 device. It may also be sent for expected changes in the medium status, if the NC is aware of them.

719 For example, if while NC-SI package is active over SMBus, the PCIe bus becomes available, this AEN
720 should be sent. Another example, if while NC-SI package is active over PCIe, the NC detects that the
721 PCIe bus is going to be disabled, it may send this AEN also.

722 In a multichannel package, the AEN, if enabled, should be sent only once per medium change event. If
723 enabled on multiple channels, the AEN may be sent on any of the channels on which this AEN is
724 enabled.

725 The media descriptors field reproduces the bit definitions defined in the Get Supported Media Response
726 (Table 11).

727

Table 15 – Medium change AEN format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved			AEN Type = 0x70
20..23	Reserved			Number of Medias supported.
24..	Media descriptors			
...	Checksum			
...	Pad			

728

729 11 Packet-Based Timing Specific to MCTP Binding

730 Table 16 presents changes in the NC-SI timing parameters relative to NC-SI Packet-Based and Op-Code
 731 Timing Parameters Table in [DSP0222](#). Parameters not listed in the table below should be taken from the
 732 table in [DSP0222](#).

733

Table 16 – NC-SI timing parameters specific to MCTP Binding

Name	Symbol	Value	Description
Normal Execution Interval	T5	50 ms, max	<p>Maximum time interval from when a controller receives a command to when it delivers a response to that command, unless otherwise specified.</p> <p>Measured from the rising edge of the first clock following the last bit of the command packet to the rising edge of the clock for the first bit of the response packet.</p> <p>Note: When T5 passed, an extension of the timeout should be allowed and taken into consideration under the following conditions:</p> <ol style="list-style-type: none"> 1. An Ethernet message or an NC-SI control message (AEN) being transmitted, 2. On a shared medium, the medium is occupied by other devices.

734

735

736
737
738

ANNEX A (informative) Notation and conventions

739 A.1 Notations

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

- 741 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets
742 starting from byte two and continuing to and including byte N. The lowest offset is on
743 the left; the highest is on the right.
- 744 • (6) Parentheses around a single number can be used in message field descriptions to
745 indicate a byte field that may be present or absent.
- 746 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range
747 may be present or absent. The lowest offset is on the left; the highest is on the right.
- 748 • PCIe Underlined, blue text is typically used to indicate a reference to a document or
749 specification called out in "Normative references" clause or to items hyperlinked within
750 the document.
- 751 • rsvd This case-insensitive abbreviation is for "reserved."
- 752 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets
753 are given as zero-based values (that is, the least significant bit [LSb] offset = 0).
- 754 • [7:5] This notation indicates a range of bit offsets. The most significant bit is on the left; the
755 least significant bit is on the right.
- 756 • 1b The lowercase "b" following a number consisting of 0s and 1s is used to indicate the
757 number is being given in binary format.
- 758 • 0x12A A leading "0x" is used to indicate a number given in hexadecimal format.

759

760
761
762
763

ANNEX B (informative)

Change log

Version	Date	Description
1.0.0	2013-08-22	
1.1.0	2015-03-21	Typos: <ul style="list-style-type: none"> • Fixed wrong message type in Table 6 Functional changes: <ul style="list-style-type: none"> • Stronger requirement on NC-SI control messages encapsulation. • Added specific timing requirements. • Added ability to send AEN on upcoming media status changes.
1.2.0	2017-08-26	<ul style="list-style-type: none"> • Updated references • Updated Contributors • Added command to enable AENs • Handled mantises • Updated list of commands supported to match NC-SI 1.1

764