



1
2
3
4

Document Identifier: DSP0261

Date: 2023-08-25

Version: 1.3.0

5 **NC-SI over MCTP Binding Specification**

6 **Supersedes: 1.2.3**

7 **Document Class: Normative**

8 **Document Status: Published**

9 **Document Language: en-US**

10 Copyright notice

11 Copyright © 2013, 2015, 2017, 2019, 2022, 2023 DMTF. All rights reserved.

12 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
13 management and interoperability. Members and non-members may reproduce DMTF specifications and
14 documents for uses consistent with this purpose, provided that correct attribution is given. As DMTF
15 specifications may be revised from time to time, the particular version and release date should always be
16 noted.

17 Implementation of certain elements of this standard or proposed standard may be subject to third party
18 patent rights, including provisional patent rights (herein “patent rights”). DMTF makes no representations
19 to users of the standard as to the existence of such rights, and is not responsible to recognize, disclose,
20 or identify any or all such third-party patent right owners or claimants, nor for any incomplete or inaccurate
21 identification or disclosure of such rights, owners, or claimants. DMTF shall have no liability to any party,
22 in any manner or circumstance, under any legal theory whatsoever, for failure to recognize, disclose, or
23 identify any such third-party patent rights, or for such party’s reliance on the standard or incorporation
24 thereof in its product, protocols or testing procedures. DMTF shall have no liability to any party
25 implementing such standard, whether such implementation is foreseeable or not, nor to any patent owner
26 or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is
27 withdrawn or modified after publication and shall be indemnified and held harmless by any party
28 implementing the standard from any and all claims of infringement by a patent owner for such
29 implementations.

30 PCI-SIG, PCIe, and the PCI HOT PLUG design mark are registered trademarks or service marks of PCI-
31 SIG. All other marks and brands are the property of their respective owners.

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 <https://www.dmtf.org/about/policies/disclosures>.

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

CONTENTS

37	Foreword	5
38	Introduction.....	6
39	Document conventions.....	6
40	Typographical conventions	6
41	ABNF usage conventions	6
42	Reserved and unassigned values.....	6
43	Byte ordering.....	6
44	Other Conventions	6
45	1 Scope	7
46	2 Normative references	7
47	3 Terms and definitions	8
48	4 Symbols and abbreviated terms.....	8
49	5 NC-SI over MCTP overview	9
50	5.1 General	9
51	5.2 NC-SI over RBT	9
52	5.3 NC-SI over MCTP.....	10
53	6 NC-SI over MCTP specific considerations	11
54	6.1 Packages and channels.....	11
55	6.2 Routing of NC-SI Pass-through traffic	13
56	6.2.1 Transmit NC-SI Pass-through traffic (MC to LAN).....	13
57	6.2.2 Receive NC-SI Pass-through traffic (LAN to MC).....	13
58	6.3 Multiple NC arbitration support	13
59	6.4 Flow control.....	14
60	6.4.1 Flow control for MCTP packets.....	14
61	6.4.2 Flow control for NC-SI over MCTP Control messages	14
62	6.4.3 Flow control for NC-SI Pass-through packets.	14
63	6.5 Interleaving of messages.....	14
64	6.5.1 Interleaving of MCTP Control and NC-SI messages	14
65	6.5.2 Interleaving of NC-SI Control and Ethernet over MCTP messages	14
66	6.6 Ordering rules for NC to MC traffic	15
67	6.7 Assembly requirements	16
68	6.8 Multiple MCTP transport bindings.....	16
69	6.8.1 Overview	16
70	6.8.2 Supported message types over different MCTP transport bindings.....	16
71	6.8.3 MCTP EID and physical address changes.	17
72	6.8.4 NC discovery flows	17
73	6.8.5 MC update flow	18
74	6.8.6 Transition between mediums	18
75	6.9 NC-SI operational state machine.....	20
76	6.9.1 Package selection.....	21
77	7 Supported NC-SI commands	21
78	8 Message types	24
79	8.1 NC-SI message type (0x02).....	24
80	8.1.1 Overview	24
81	8.1.2 Encapsulation	24
82	8.2 Ethernet message type (0x03).....	26
83	8.2.1 Overview	26
84	8.2.2 Encapsulation	26
85	8.3 Version.....	27
86	9 NC-SI commands and AENs specific to MCTP transport.....	28
87	9.1 Overview	28

88 9.2 Get Supported Media Command (0x54)..... 28

89 9.3 Get Supported Media Response (0xD4) 28

90 9.4 Transport Specific AENs Enable (0x55) 29

91 9.5 Transport Specific AENs Enable Response (0xD5)..... 30

92 9.6 Medium Change AEN 30

93 10 Packet-Based Timing Specific to MCTP Binding 31

94 ANNEX A (informative) Notation and conventions 32

95 ANNEX B (informative) Change log..... 33

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 – Ordering rules 15

107 Table 3 – NC-SI operating state relevance to MCTP binding..... 21

108 Table 4 – MCTP-specific NC-SI commands requirements 22

109 Table 5 – NC-SI messages encapsulation..... 25

110 Table 6 – MCTP Transport Header fields 25

111 Table 7 – MCTP Specific Message Header field 25

112 Table 8 – Ethernet messages encapsulation..... 26

113 Table 9 – MCTP Transport Header fields 27

114 Table 10 – MCTP Specific Message Header field 27

115 Table 11 – Get Supported Media Command Packet Format..... 28

116 Table 12 – Get Supported Media Response Packet Format 28

117 Table 13 – Get Supported Media Response media descriptors format..... 29

118 Table 14 – Transport Specific AENs Enable Command Packet Format 29

119 Table 15 – Transport Specific AENs enable field format 30

120 Table 16 – Transport Specific AENs Enable Response Packet Format..... 30

121 Table 17 – Medium change AEN format..... 31

122 Table 18 – NC-SI Timing Parameters Specific to MCTP Binding 31

123

Foreword

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

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

128 This version supersedes version 1.2.3. For a list of changes, see the change log in ANNEX B

129 **Acknowledgments**

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

131 **Editor:**

- 132 • Eliel Louzoun – Intel Corporation

133 **Contributors:**

- 134 • Alan Berenbaum – SMSC
- 135 • Patrick Caporale – Lenovo
- 136 • Philip Chidester – Dell Technologies
- 137 • Kelly Couch – Intel Corporation
- 138 • Benzi Friedman – Intel Corporation
- 139 • Yuval Itkin – NVIDIA Corporation
- 140 • Ira Kalman – Intel Corporation
- 141 • Ed Klodnicki – IBM
- 142 • Joe Kozlowski – Dell Technologies
- 143 • Patrick Kutch – Intel Corporation
- 144 • John Leung – Intel Corporation
- 145 • Patrick Schoeller – Intel Corporation
- 146 • Hemal Shah – Broadcom Inc.
- 147 • Tom Slaight – Intel Corporation
- 148 • Bob Stevens – Dell Technologies

149

Introduction

150 The *NC-SI over MCTP Binding Specification* defines new MCTP messages used to convey NC-SI Control
151 packets and Ethernet traffic over MCTP to allow NC-SI Pass-through traffic over MCTP. This specification
152 is based on the [DSP0222 1.2](#) specification and uses the same NC-SI Control packet definitions.

153 Document conventions

154 Typographical conventions

155 The following typographical conventions are used in this document:

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

161 ABNF usage conventions

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

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

166 Reserved and unassigned values

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

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

171 Byte ordering

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

174 Other Conventions

175 See ANNEX A for other conventions.
176

177 1 Scope

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

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

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

187 2 Normative references

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

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

195 DMTF DSP0004, *CIM Infrastructure Specification 3.0*,
196 https://www.dmtf.org/standards/published_documents/DSP0004_3.0.1.pdf

197 DMTF DSP0222, *Network Controller Sideband Interface (NC-SI) Specification 1.2*
198 https://www.dmtf.org/sites/default/files/standards/documents/DSP0222_1.2.0.pdf

199 DMTF DSP0223, *Generic Operations 1.0*,
200 https://www.dmtf.org/standards/published_documents/DSP0223_1.0.pdf

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

203 DMTF DSP0237, *Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding*
204 *Specification 1.2*
205 https://www.dmtf.org/standards/published_documents/DSP0237_1.2.pdf

206 DMTF DSP0238, *Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding*
207 *Specification 1.2*
208 https://www.dmtf.org/standards/published_documents/DSP0238_1.2.pdf

209 DMTF DSP0239, *Management Component Transport Protocol (MCTP) IDs and Codes 1.7*
210 https://www.dmtf.org/standards/published_documents/DSP0239_1.7.pdf

211 DMTF DSP1001, *Management Profile Specification Usage Guide 1.2*,
212 https://www.dmtf.org/standards/published_documents/DSP1001_1.2.pdf

213 ACPI, *Advanced Configuration and Power Interface Specification Revision 4.0a*, April 5, 2010
214 <https://web.archive.org/web/20190416120448/http://www.acpi.info/DOWNLOADS/ACPIspec40a.pdf>

215 IETF RFC5234, *ABNF: Augmented BNF for Syntax Specifications, January 2008*,
216 <https://datatracker.ietf.org/doc/html/rfc5234>

217 ISO/IEC Directives, Part 2, *Principles and rules for the structure and drafting of ISO and IEC documents*,
218 <https://www.iso.org/sites/directives/current/part2/index.xhtml>

219 **3 Terms and definitions**

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

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

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

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

233 The terms defined in [DSP0004](#), [DSP0223](#), [DSP0236](#) and [DSP1001](#) apply to this document. The following
234 additional terms are used in this document.

235 **3.1**

236 **System Power States**

237 **S0 and Sx**

238 S0 represents an active system

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

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

241 **4 Symbols and abbreviated terms**

242 The abbreviations defined in [DSP0004](#), [DSP0223](#), [DSP0236](#), and [DSP1001](#) apply to this document. The
243 following additional abbreviations are used in this document.

244 **4.1**

245 **ACPI**

246 Advanced Configuration and Power Interface

247 **4.2**

248 **IANA**

249 Internet Assigned Numbers Authority

250 **4.3**

251 **FCS**

252 Frame Check Sequence

253 **4.4**

254 **MCTP**

255 Management Component Transport Protocol

256 **4.5**
257 **MC**
258 Management Controller

259 **4.6**
260 **NC**
261 Network Controller

262 **4.7**
263 **NC-SI**
264 Network Controller Sideband Interface

265 **4.8**
266 **RID**
267 PCIe Requester ID (Bus/Device/Function).

268 **5 NC-SI over MCTP overview**

269 **5.1 General**

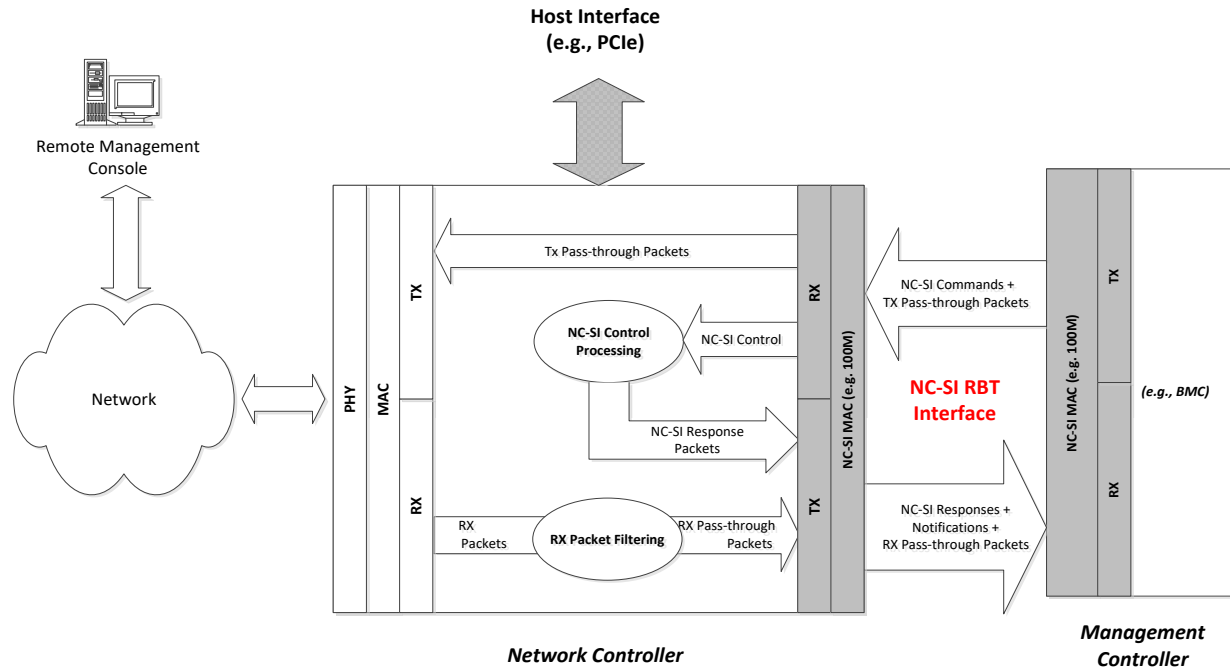
270 NC-SI over MCTP is based on DSP0222 ([NC-SI](#)). The *NC-SI over MCTP Binding Specification* replaces
271 the RBT Protocol with a definition of NC-SI communications using MCTP. The MCTP Transport Bindings
272 are defined in other companion specifications such as *MCTP SMBus Binding Specification* ([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 are
275 relevant to this specification.

276 **5.2 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 based on a specific EtherType (0x88F8).

293

294

- A control protocol defining a set of commands allowing an MC to configure and monitor Network Controllers and their Pass-through channels for MC to network communication. The command set functionality can be extended using OEM commands.

295

296

297 5.3 NC-SI over MCTP

298

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

299

300

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

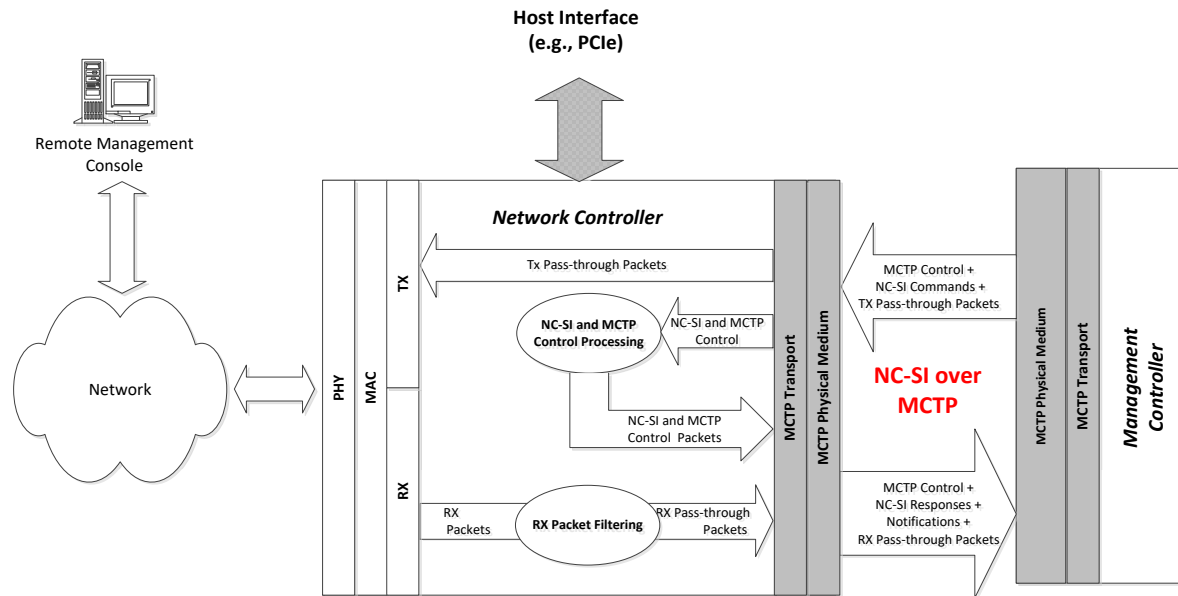
301

302

303

304

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 8.1.2. The
 316 encapsulation of NC-SI Pass-through traffic in MCTP messages is described in subclause 8.2.2.

317 6 NC-SI over MCTP specific considerations

318 6.1 Packages and channels

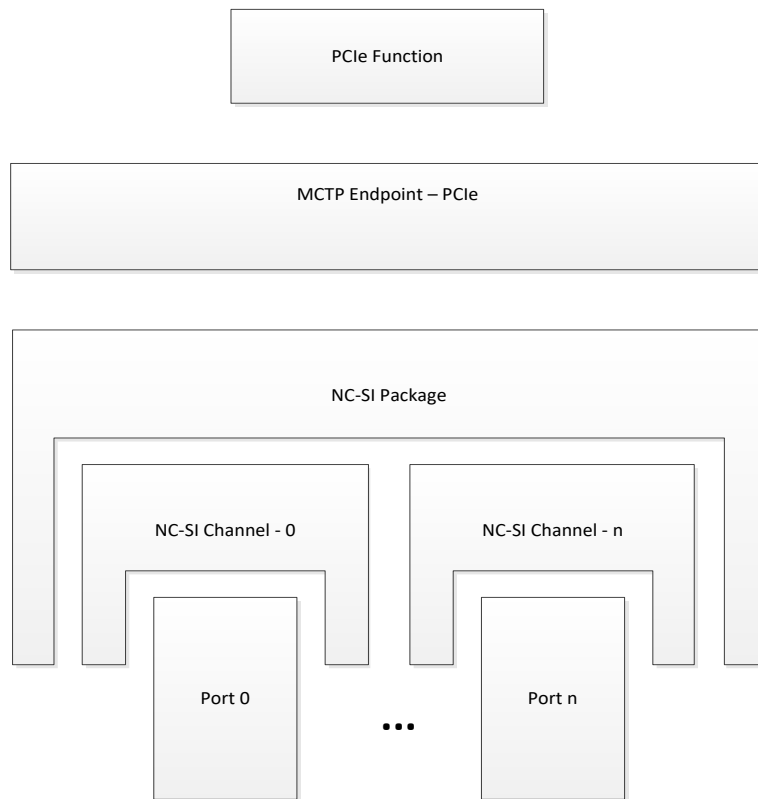
319 The NC-SI specification defines different topologies using the concepts of channels and packages. A
 320 channel is associated with a network port and a package is usually associated with a physical device that

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

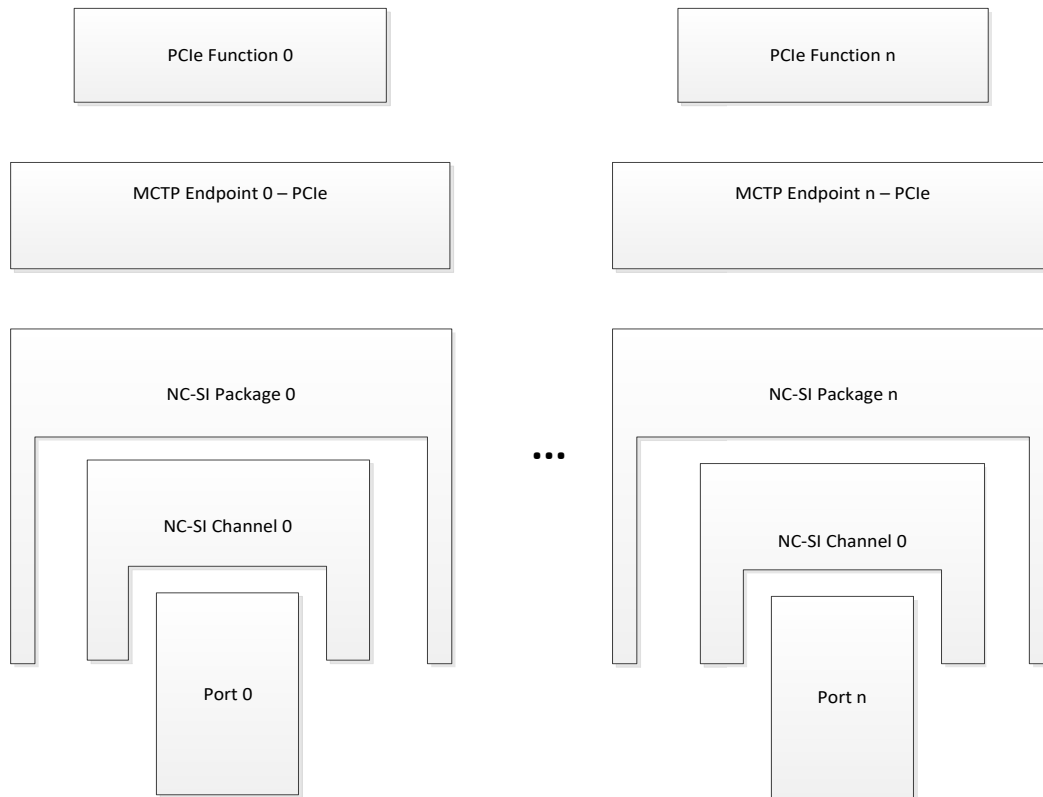


333

334

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

335



336

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

337

338 Multiple MCTP transport bindings handling is described in subclause 6.8.

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

340 **6.2 Routing of NC-SI Pass-through traffic**

341 **6.2.1 Transmit NC-SI Pass-through traffic (MC to LAN)**

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

349 **6.2.2 Receive NC-SI Pass-through traffic (LAN to MC)**

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

351 **6.3 Multiple NC 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 **6.4 Flow control**

359 **6.4.1 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 **6.4.2 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 **6.4.3 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 **6.5 Interleaving of messages**

379 **6.5.1 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 **6.5.2 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 **6.6 Ordering rules for NC to MC traffic**

394 Table 2 defines which type of messages should pass other types of packets to avoid deadlocks. The
 395 decisions are done at a message level. Interleaving within messages is defined in the previous sections.
 396 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 **Table 2 – Ordering rules**

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 • **A** This row relates only to the precedence of MCTP base control messages
 406 over NC-SI and Ethernet messages and not over other MCTP message
 407 types.
- 408 • **A1** This situation will occur only in NCs accepting multiple outstanding
 409 MCTP control commands.
- 410 • **B2** This situation will occur only in NCs accepting multiple outstanding NC-
 411 SI commands.
- 412 • **A3, B3** An NC-SI AEN might be blocked if the channel is disabled or the
 413 package deselected. Thus, it should not block MCTP Control or NC-SI
 414 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-
 417 down event followed by a link-up event, the AEN on the link-up event
 418 must not 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 **6.7 Assembly requirements**

422 According to the interleaving requirements described in subclause 6.5, 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 8.2. The maximum NC-SI packet size is defined in subclause 8.1.

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

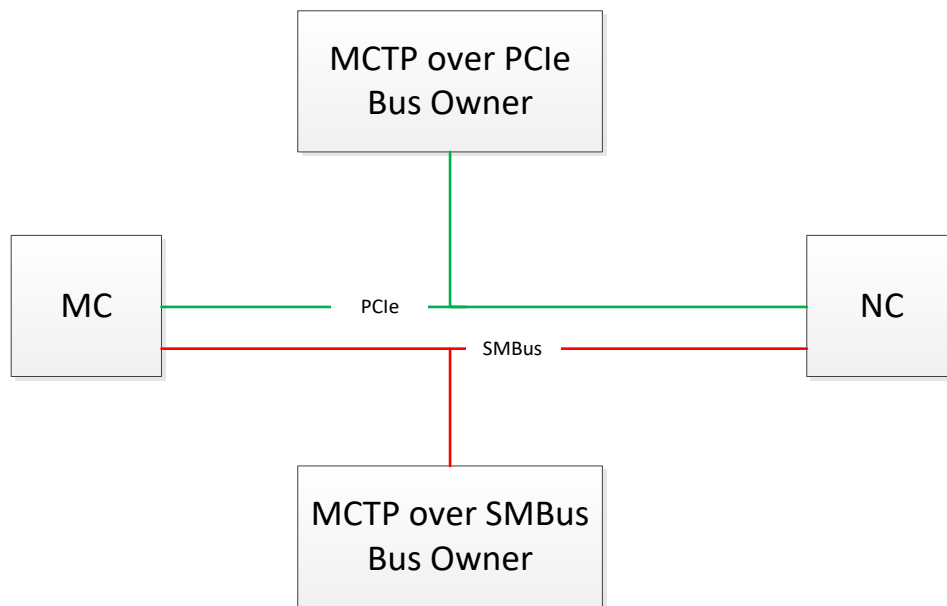
426 **6.8 Multiple MCTP transport bindings**

427 **6.8.1 Overview**

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

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

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



437

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

439 **6.8.2 Supported message types over different MCTP transport bindings**

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

445

446 **6.8.3 MCTP EID and physical address changes.**

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

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

456 In order to avoid breakup of network connections, and in order to avoid the need to reconfigure the NC,
457 the NC-SI connection should be kept alive during the transition. The MC is responsible for the
458 reconnection of the channel in case of address mapping changes. The next clause describes possible
459 flows that may be used to ease the re-discovery of an NC whose address has changed. A flow by which
460 the MC can expose a change of its own address to the NC(s) is described in subclause 6.8.5.

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

463 **6.8.4 NC discovery flows**

464 **6.8.4.1 General**

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

466 **6.8.4.2 Full discovery**

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

469 The following flow may be used:

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

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

481 **6.8.4.3 UUID based discovery**

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

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

485 The following flow may be used:

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

501 **6.8.4.4 NC-SI based discovery**

502 The NC must support the “Get Supported Media” NC-SI command as defined in clause 9.2 to use this
503 method.

504 The following flow may be used.

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

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

516 **6.8.5 MC update flow**

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

519 **6.8.6 Transition between mediums**

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

- 523 A transition of an NC-SI from one medium to another is achieved when the NC is deselected on the first
524 medium and selected on the second medium as described in subclause 6.9.
- 525 The NC may notify the MC about the state of a medium using an AEN.
- 526 1) Potential loss of a medium prior to losing the medium
- 527 2) Availability of a new medium
- 528 Alternatively, the MC may be aware of the medium change independently, for example, by detecting its
529 own PCIe bus became active, by interaction with the BIOS, and so on.
- 530 The MC may initiate the transition by using MCTP Control and NC-SI Control messages as described
531 below.
- 532 A transition can be between mediums with different levels of support of Ethernet MCTP messages.
- 533 When an NC transitions from a medium on which Ethernet messages were supported to a second
534 medium on which Ethernet messages are not supported, the NC should stop sending and receiving
535 Ethernet messages on the first medium after the NC-SI channel had been deselected on the first medium.
- 536 The MC may transition back later to the first medium for communicating Ethernet messages. If the MC
537 transitions back to the first medium supporting Ethernet messages, it may resume communications of
538 Ethernet messages based on the previous configuration. If the configuration was lost during the
539 transitions, the NC shall return to the NC-SI Initial State (as described in section 6.2.4 of [NC-SI](#)).
- 540 Even if NC-SI Pass-through traffic (Ethernet messages) is supported over multiple mediums, Pass-
541 through traffic shall not be transitioned to a new medium before the connection between the MC and the
542 NC is re-established on the new medium. The NC shall support the following flows to initiate a transition
543 to the new medium:
- 544 • If the current medium is still active (for example when moving from SMBus to PCIe to achieve
545 better throughput), the NC shall keep its Pass-through traffic on the original medium (both MC to
546 network and network to MC). The NC shall also send outstanding NC-SI responses on the
547 original medium.
548 NOTE The MC can stop the traffic from the NC on the current medium by sending “Disable
549 Channel” and “Disable Channel Network TX” NC-SI commands to all the channels before the
550 transition. In this case, it can send “Enable Channel” and “Enable Channel Network TX” NC-SI
551 commands to all active channels on the new medium, to allow the traffic to resume.
 - 552 • If the current medium is inactive (for example, when moving from PCIe to SMBus due to a power
553 transition), the NC shall stop transferring Ethernet messages. If a packet is being transmitted by
554 the NC when the original medium becomes unavailable, the NC shall not continue the
555 transmission of the packet and the packet might be lost. Outstanding NC-SI responses may be
556 discarded by the NC.
 - 557 • When any NC-SI command is received from the MC on the new medium (apart from “*Deselect*
558 *Package*”), the NC shall move to “Selected” state on the new medium (see subclause 6.9).
 - 559 ○ An NC-SI Rx Pass-through message to the MC on the current medium shall be
560 completed by the NC on the current medium and only after that shall the NC send the
561 NC-SI response to the MC on the new medium.
 - 562 ○ The next Pass-through message sent to the MC after a successful response to the NC-SI
563 command shall be sent on the new medium.

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

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

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

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

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

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

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

578 • If the current medium is still active (for example when moving from SMBus to PCIe to achieve
579 better throughput), the MC can keep its traffic on the original medium until it discovers the NC by
580 using one of the flows described in subclause 6.8.4. If the current medium is inactive (for
581 example, when moving from PCIe to SMBus due to a power transition), the MC will stop
582 transferring Ethernet messages with NC until discovery of the NC.

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

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

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

594 AENs may require a behavior different from Ethernet traffic when transitioning between media. For
595 example, when a PCIe link is disabled, the NC may need to indicate this condition over SMBus before the
596 MC transitions the Ethernet traffic from the PCIe medium to the SMBus medium. Also, the enablement of
597 AENs on one medium doesn't mean it is applicable for another medium. Thus, AENs shall be enabled on
598 each medium independently, i.e., the AEN Enable (0x08) and Transport Specific AENs Enable (0x55)
599 commands are medium-specific. An AEN enabled over PCIe may be disabled over SMBus and vice
600 versa.

601 **6.9 NC-SI operational state machine**

602 [DSP0222](#) § "Operational behaviors" defines a set of operational states an NC may be in. Not all those
603 states are relevant to MCTP binding. The table below lists which of the states are relevant to MCTP
604 binding:

605

606

Table 3 – NC-SI operating state relevance to MCTP binding

State	Applies to	Applicable to This Binding?
Interface Power Down	Package	Yes
Interface Power Up	Package	Yes
Package Selected (also referred to as the Selected state)	Package	Yes
Package Deselected (also referred to as the Deselected state)	Package	Yes
Hardware Arbitration Enabled	Package	No
Hardware Arbitration Disabled	Package	No
Package Ready	Package	Yes
Package Not Ready	Package	Yes
Channel Ready	Channel	Yes
Channel Not Ready	Channel	Yes
Initial State	Channel	Yes
Channel Enabled	Channel	Yes
Channel Disabled	Channel	Yes

607

608 6.9.1 Package selection

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

611 As in [DSP0222](#), a package is selected by reception of a “Select Package” on the MCTP medium or any
 612 other command except “Deselect Package”.

613 A package is deselected on a specific MCTP medium by reception of a “Deselect Package” command,
 614 selection of the package on another medium or if the physical medium on which it operate becomes
 615 unavailable. If the packet is deselected by an NC-SI command it should move to the deselected state only
 616 after sending a response to the command.

617 A package is allowed to send Ethernet messages or NC-SI Control messages on an MCTP medium only
 618 if in the “Selected” state on that medium.

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

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

625 7 Supported NC-SI commands

626 The supported NC-SI commands when bound to MCTP is a subset of the commands in [DSP0222](#)
 627 specification. The subset of supported commands varies according to the supported messages as
 628 indicated in the response to the Get Message Type Support MCTP Control command. If only the NC-SI
 629 Control message type is supported, the commands related to the Pass-through traffic control are not
 630 supported. If both the NC-SI Control and Ethernet message types are supported, these commands are
 631 supported. Table 4 lists the supported commands according to the supported message types wherever
 632 the support differs from the support listed in [DSP0222](#).

633 Optional commands may have different implementation over different media.

634 Note that some commands are not applicable for MCTP binding and are listed here only for
 635 completeness. These commands are marked as “N/A”.

636

637 **Table 4 – MCTP-specific NC-SI commands requirements**

Command Type	Command Name	Description	Response Type	Command Support Requirement: NC-SI Control Messages Only	Command Support Requirement: NC-SI Control and Ethernet Messages
0x01	Select Package	Used to explicitly select a controller package to transmit packets through the NC-SI interface	0x81	O ³	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
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
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

Command Type	Command Name	Description	Response Type	Command Support Requirement: NC-SI Control Messages Only	Command Support Requirement: NC-SI Control and Ethernet Messages
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
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	M ²	M
0x1A	Get NC-SI Pass-through Statistics	Used to request NC-SI Pass-through packet statistics	0x9A	N/A	O
0x52	Get Package UUID	Returns a universally unique identifier (UUID) for the package	0xD2	O	O
0x53	Reserved	Used by DSP0222 for RBT binding	0xD3	N/A	N/A
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
0x56	Reserved	Used by DSP0222 for RBT binding	0xD6	N/A	N/A
0x57	Reserved	Used by DSP0222 for RBT binding	0xD7	N/A	N/A
0x60	Reserved	Used by DSP0222 for RBT binding	0xE0	N/A	N/A
0x61	Reserved	Used by DSP0222 for RBT binding	0xE1	N/A	N/A

Command Type	Command Name	Description	Response Type	Command Support Requirement: NC-SI Control Messages Only	Command Support Requirement: NC-SI Control and Ethernet Messages
0x62	Reserved	Used by DSP0222 for RBT binding	0xE2	N/A	N/A
Key: M = Mandatory (required) O = Optional C = Conditional (see command description) N/A = Not applicable					
1. The only part of the response that is relevant is the AEN control support field. 2. 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. 3. 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.					

638 **8 Message types**

639 **8.1 NC-SI message type (0x02)**

640 **8.1.1 Overview**

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

643 This message type shall be supported in any device compliant with this specification

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

646 **8.1.2 Encapsulation**

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

650 NC-SI control packets communicated over MCTP do not follow the Ethernet frame encapsulation defined
 651 in DSP0222 for NC-SI over RMI Based Transport (RBT) transport binding. NC-SI control packets over
 652 MCTP shall not include Ethernet frame header, Ethernet packet pad, and Ethernet Frame Check
 653 Sequence (FCS). Instead, the encapsulation described in Table 5 shall be used to encapsulate NC-SI
 654 control messages.

655 NOTE The Control packets frames in DSP0222 use a DA, SA, and EtherType MAC header. The DA and SA part do
 656 not contain any useful data and the EtherType is used to differentiate between Control packets and Ethernet traffic. In
 657 NC-SI over MCTP, this Ethernet framing is not used, as the differentiation is achieved through usage of different
 658 message types.

659

660

Table 5 – 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							
00..03	RSVD				Header Version				Destination Endpoint ID				Source Endpoint ID				S	E	Pkt	T	Message	O	O	seq #	O	Tag	M	M			
04..07	IC	Message Type				MC ID				Header Revision				Reserved				0													
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																														

661

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

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

665

Table 6 – 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.

666

667

Table 7 – 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).

668

669

670 **8.2 Ethernet message type (0x03)**

671 **8.2.1 Overview**

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

674 This message type should be supported in any device compliant with this specification that supports pass
675 through traffic.

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

679 **8.2.2 Encapsulation**

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

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

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

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

689 **Table 8 – 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								S	E	Pkt	T	Message			
																								M	O	seq #	O	Tag				
04..07	IC	Message Type							Destination Address 5:3																							
	0	0x03																							Source Address 5							
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)																															

690
691 The following tables describe the value for the various fields of the message whose description differs
692 from the description in the MCTP or NC-SI specification.

693

Table 9 – 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.

694

695

Table 10 – 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).

696

697 8.3 Version

698 The versions that shall be reported for NC-SI control or Ethernet message types in the Get MCTP Version
699 Support response are as follow:

- 700 • The Version Number Entry 1 field shall be used to indicate backward compatibility with Version
701 1.0 as:
702 1.0 [Major version 1, minor version 0, any update version, no alpha]
703 This is reported using the encoding as: 0xF1F0FF00
704
- 705 • The Version Number Entry 2 field shall be used to indicate backward compatibility with Version
706 1.1 as:
707 1.1 [Major version 1, minor version 1, any update version, no alpha]
708 This is reported using the encoding as: 0xF1F1FF00
709
- 710 • The Version Number Entry 3 field shall be used to indicate backward compatibility with Version
711 1.2 as: 1.2 [Major version 1, minor version 2, any update version, no alpha]
712 This is reported using the encoding as: 0xF1F2FF00
713
- 714 • The version for this specification shall be reported in Version Number Entry 4 as:
715 1.3.0 [Major version 1, minor version 3, update version 0, no alpha]
716 This is reported using the encoding as: 0xF1F3F000
717
718

719 **9 NC-SI commands and AENs specific to MCTP transport**

720 **9.1 Overview**

721 The following commands and AEN may be implemented as part of this specification to allow an
722 implementation of the discovery flow described in clause 6.8.4.4.

723 **9.2 Get Supported Media Command (0x54)**

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

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

729 Table 11 illustrates the packet format of the Get Supported Media command.

730 **Table 11 – 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			

731 **9.3 Get Supported Media Response (0xD4)**

732 In the absence of any error, the package shall process and respond to the Get Supported Media
733 command by sending the response packet and payload shown in Table 12.

734 **Table 12 – Get Supported Media Response Packet Format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Reserved			Number of medias supported
24..	Media descriptors as described in Table 13. The number of media descriptors is according to the Number of medias supported field value.			
...	Checksum			
...	Pad			

735

736

Table 13 – 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	Bit 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. Bits 6:1: Reserved Bit 7: Status 0: Medium is not currently available for NC-SI. 1: Medium is currently available for NC-SI.
4	Physical Address Size. Should be 0x0 if Physical Medium Identifier is NC-SI over RBT, otherwise shall be according to MCTP binding specifications.
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.

737 The fields of the media descriptors are considered separate fields in the response; hence, Byte 0 of the
 738 descriptor is transmitted first. The Physical address field, for Physical address sizes greater than one, is a
 739 multibyte field and is transmitted most significant byte first.

740 **9.4 Transport Specific AENs Enable (0x55)**

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

747 Enablement of AENs on one medium is not applicable to other media.

748 Table 14 illustrates the packet format of the Enable Transport Specific AENs command.

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

750 **Table 14 – 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			

751

752

Table 15 – 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

753

754 9.5 Transport Specific AENs Enable Response (0xD5)

755 In the absence of any error, the package shall process and respond to the Transport Specific AENs
756 Enable command by sending the response packet and payload shown in Table 16.

757

Table 16 – Transport Specific AENs Enable Response Packet Format

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

758

759 9.6 Medium Change AEN

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

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

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

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

771 The media descriptors field reproduces the bit definitions defined in the Get Supported Media Response
772 (Table 13).

773

Table 17 – Medium change AEN format

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

774 **10 Packet-Based Timing Specific to MCTP Binding**

775 Table 18 presents changes in the NC-SI timing parameters relative to NC-SI Packet-Based and Op-Code
 776 Timing Parameters Table in [DSP0222](#). Parameters not listed in the table below should be taken from the
 777 table in [DSP0222](#).

778

Table 18 – NC-SI Timing Parameters Specific to MCTP Binding

Name	Symbol	Value	Description
Normal Execution Interval	T5	50 ms, max	Maximum time interval from when a controller receives a command to when it delivers a response to that command, unless otherwise specified. 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. Note: When T5 passed, an extension of the timeout should be allowed and taken into consideration under the following conditions: <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.

779

ANNEX A (informative)

Notation and conventions

780
781
782
783

784 Notations

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

- 786 • 2:N In field descriptions, this will typically be used to represent a range of byte offsets
787 starting from byte two and continuing to and including byte N. The lowest offset is on
788 the left; the highest is on the right.
- 789 • (6) Parentheses around a single number can be used in message field descriptions to
790 indicate a byte field that may be present or absent.
- 791 • (3:6) Parentheses around a field consisting of a range of bytes indicates the entire range
792 may be present or absent. The lowest offset is on the left; the highest is on the right.
- 793 • PCIe Underlined, blue text is typically used to indicate a reference to a document or
794 specification called out in "Normative references" clause or to items hyperlinked within
795 the document.
- 796 • rsvd This case-insensitive abbreviation is for "reserved."
- 797 • [4] Square brackets around a number are typically used to indicate a bit offset. Bit offsets
798 are given as zero-based values (that is, the least significant bit [LSb] offset = 0).
- 799 • [7:5] This notation indicates a range of bit offsets. The most significant bit is on the left; the
800 least significant bit is on the right.
- 801 • 1b The lowercase "b" following a number consisting of 0s and 1s is used to indicate the
802 number is being given in binary format.
- 803 • 0x12A A leading "0x" is used to indicate a number given in hexadecimal format.

804
805
806
807
808

809
810
811
812

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 8 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	Updated references Updated Contributors Added command to enable AENs Handled mantises Updated list of commands supported to match NC-SI 1.1
1.2.1	2018-08-23	Added reason code and response code to response format
1.2.2	2019-09-24	Fixed reported versions
1.2.3	2021-05-14	Updated to comply with ISO guidelines.
1.3.0	2023-08-25	Updated references versions. Changed command table to delta table Updated reported versions (for 1.2.3 and 1.3.0) AENs are medium specific Added table of states relevant to MCTP binding Unified the version reporting section for both types Added byte ordering for media descriptor

813